

HP OpenView Storage Mirroring

User's Guide



© Copyright 2005 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Microsoft®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

HP OpenView Storage Mirroring User's Guide

Sixth Edition (June 2005)

Part Number: T2558-90010

NSI and Double-Take are registered trademarks of NSI Software, Inc. All other products are trademarks of their respective companies.
(c) 1996-2005 NSI Software, Inc.

Table of Contents

Preface	i
Intended Audience	i
Organization of This Guide	i
Conventions Used in This Guide	i
Online References	ii
Online help	ii
Online manual	ii
Chapter 1: Introduction to Storage Mirroring	1-1
Storage Mirroring Features	1-2
High availability and disaster recovery management features	1-2
Resource utilization features	1-2
Improved backup benefits	1-2
Other features and benefits	1-2
Storage Mirroring Operations	1-3
Mirroring	1-3
Replication	1-4
Failure monitoring/failover	1-5
Restoration	1-6
Sample Configurations	1-7
One-to-One, Active/Standby	1-7
One-to-One, Active/Active	1-8
Many-to-One	1-9
One-to-Many	1-10
Chained	1-11
Single Machine	1-12
Chapter 2: Management Console	2-1
Starting the Management Console	2-2
Understanding the Display	2-3
Management Console toolbar	2-5
Logging On and Off of Storage Mirroring	2-6
Stopping the Service	2-6
Storage Mirroring Ports	2-7
Management Console ports	2-7
Storage Mirroring server ports	2-8
Managing the Management Console Tree	2-9
Groups	2-9
Creating groups	2-9
Removing groups	2-9
Servers	2-9
Moving servers	2-10
Insert a server	2-10
Remove a server	2-10
Hide a server	2-11
Unhide a server	2-11
Sharing group and server configuration	2-11

Storage Mirroring Workspaces	2-12
Saving a workspace	2-12
Opening a workspace	2-12
Getting Help for the Management Console	2-12
Exiting the Management Console	2-12
Chapter 3: Text Clients	3-1
The Text Client	3-2
Starting the Text Client	3-2
Using the Text Client	3-2
The Command Line Client	3-3
Starting and using the Command Line Client	3-3
Logging on and off of Storage Mirroring	3-4
Stopping the Service	3-5
Specifying Source and Target Functionality	3-6
Storage Mirroring Ports	3-7
Text client ports	3-7
Storage Mirroring server ports	3-9
Getting Help in the Text Clients	3-10
Exiting the Text Clients	3-10
Chapter 4: Failover Control Center	4-1
Starting the Failover Control Center	4-2
Understanding the Display	4-3
Identifying what the target is monitoring	4-3
Configuring failover operations	4-3
Monitoring failover	4-4
Communication ports	4-6
Chapter 5: Replication Sets	5-1
Replication Set Rules	5-2
Replication Set Operations	5-3
How Replication Sets Work	5-4
Multi-Client Access	5-5
Replication Capabilities	5-6
Managing Replication Sets Through the Management Console	5-7
Creating a replication set	5-7
Creating a replication set using wild cards and direct rule entry	5-9
Calculating replication set size	5-10
Modifying a replication set	5-11
Renaming and copying replication sets	5-11
Deleting a replication set	5-11
Managing Replication Sets Through the Text Client	5-12
Creating a replication set	5-12
Calculating replication set size	5-14
Modifying a replication set	5-16
Deleting a replication set	5-19

Chapter 6: Connections	6-1
Connection Options	6-2
Throughput Diagnostic Utility	6-4
Connection IDs	6-4
Connection Resources	6-4
Compression	6-4
How Connections Are Established	6-5
Queuing	6-6
Auto-Disconnect and Auto-Reconnect	6-7
Managing Connections Through the Management Console	6-8
Establishing a new connection	6-8
Disconnecting a connection	6-9
Monitoring a connection	6-10
Connection statistics	6-10
Connection display	6-11
Configuring compression	6-12
Configuring queuing	6-13
Configuring auto-reconnect	6-15
Managing Connections Through the Text Client	6-16
Establishing a new connection	6-16
Connecting to the Throughput Diagnostics Utility	6-18
Disconnecting a connection	6-19
Monitoring a connection	6-20
Storing the connection ID in a variable	6-21
Configuring compression	6-22
Configuring queuing	6-23
Configuring auto-reconnect	6-24
Chapter 7: Mirroring	7-1
Mirroring Options	7-1
Stopping and Pausing Mirroring	7-3
Orphan Files	7-3
Auto-Remirror	7-4
How Mirroring Works	7-5
Managing Mirroring Using the Management Console	7-6
Initiating a mirror upon connection	7-6
Controlling a mirror manually	7-7
Monitoring a mirror	7-8
Configuring orphan files	7-9
Setting auto-remirror	7-10
Managing Mirroring Using the Text Client	7-11
Initiating a mirror on connection	7-11
Controlling a mirror manually	7-13
Monitoring a mirror	7-15
Configuring orphan files	7-16
Setting auto-remirror	7-18

Chapter 8: Replication	8-1
Replication Options	8-1
How Replication Works	8-1
Recovery From Network Losses	8-2
Windows Permissions	8-2
Task Command Processing	8-2
Managing Replication Through the Management Console	8-3
Initiating replication upon connection	8-3
Controlling replication manually	8-4
Monitoring replication	8-5
Replicating permissions by name	8-6
Processing task commands	8-7
Managing Replication Through the Text Client	8-8
Initiating replication on connection	8-8
Controlling replication manually	8-10
Monitoring replication	8-11
Replicating permissions by name	8-12
Processing task commands	8-13
Chapter 9: Verification	9-1
Verification Options	9-2
Verification Log	9-4
How Verification Works	9-6
Managing Verification With the Management Console	9-7
Configuring verification on connection	9-7
Initiating verification manually	9-8
Verification log	9-9
Monitoring verification	9-10
Managing Verification Using the Text Client	9-11
Initiating verification	9-11
Setting the verification log attributes	9-13
Monitoring verification	9-14
Chapter 10: Data Transmission	10-1
Bandwidth Limitations	10-2
Real-Time and Manual Transmission	10-2
Scheduled Transmission	10-2
Scheduled Transmission Options	10-3
Setting Transmission Criteria Using the Management Console	10-5
Establishing transmission criteria when configuring a connection	10-5
Setting bandwidth limitations	10-6
Setting transmission start criteria	10-7
Setting transmission stop criteria	10-8
Setting a transmission window	10-9
Controlling transmission manually	10-10

Setting Transmission Criteria Using the Text Client	10-11
Configuring transmission criteria before establishing a connection	10-11
Clearing all transmission criteria	10-11
Enabling and disabling transmission criteria	10-11
Setting bandwidth limitations	10-12
Setting transmission start criteria	10-13
Setting transmission stop criteria	10-14
Setting a transmission window	10-15
Transmission schedule examples	10-16
Controlling transmission manually	10-17
Pausing Storage Mirroring Execution on the Target	10-18
Pausing through the Management Console	10-18
Pausing through the text clients	10-18
Resuming through the Management Console	10-19
Resuming through the text clients	10-19
Chapter 11: Failover	11-1
The Failover Process	11-2
The Failback Process	11-3
Failover Clients	11-4
Failover Terminology	11-5
How Failover Works	11-6
How Failback Works	11-8
How Failure Detection Works	11-9
Client Connectivity	11-9
Operating System Differences	11-10
Windows 200x	11-10
Windows NT 4.0	11-11
Failover Hardware Configurations	11-12
Single NIC on the source and target	11-12
Multiple NICs on the source and single NIC on the target	11-13
Multiple NICs on the source and target	11-14
Failover Options	11-15
Managing Failover Through the Failover Control Center	11-17
Configuring monitoring and failover	11-17
Monitoring failover	11-21
Initiating failback	11-24
Failover IP address placeholders	11-25
Updating the number of IP address placeholders	11-25
Communication ports	11-26
Manual share updates	11-26
Failover Control Center window refresh rate	11-26
Testing failover	11-27
Failover manual intervention	11-27
Managing Failover Through the Text Client	11-28
Configuring monitoring and failover	11-28
Testing failover	11-37
Initiating failback	11-38

Chapter 12: Restoration	12-1
How the Restoration Process Works	12-2
Restoration Options	12-3
Restoring Storage Mirroring Data Through the Management Console	12-5
Restoring Storage Mirroring Data Through the Text Client	12-7
Automating Restore	12-9
 Chapter 13: Monitoring Tools	 13-1
Management Console Tools	13-2
Management Console statistics	13-2
Customizing the statistics	13-4
Changing the order of the statistics	13-5
Saving statistics modifications	13-5
Reverting to the last configuration	13-5
Resetting the default configuration	13-5
At-a-Glance monitoring	13-6
Customizing when the background color and icons change	13-9
Management Console message window	13-10
Text Client Tools	13-12
Text Client main window	13-12
DTCL status command	13-13
Operating System Dependent Tools	13-14
Windows Event Viewer	13-14
Event descriptions	13-15
Windows Performance Monitor	13-27
Performance Monitor statistics	13-27
Operating System Independent Tools	13-30
DTStat	13-30
DTStat log file	13-30
Running DTStat	13-33
DTStat statistics	13-35
LogViewer	13-39
LogViewer log file	13-40
LogViewer Messages	13-44
SNMP	13-48
Configuring SNMP	13-48
SNMP statistics	13-48
SNMP traps	13-50
 Chapter 14: Advanced Settings	 14-1
Changing the Settings Through the Management Console	14-1
General options	14-2
Setup options	14-4
Network options	14-5
Queue options	14-6
Source options	14-7
Target options	14-9
Database options	14-10
Logging options	14-11
Notification options	14-12

Changing the Settings Through the Text Client	14-14
Storage Mirroring Text Client program settings	14-15
Chapter 15: Security	15-1
Security Access Levels	15-1
Security Advantages and Considerations	15-1
How Storage Mirroring Security Works	15-2
Clearing Maintained Security Credentials	15-3
Storage Mirroring Service	15-4
Storage Mirroring Groups	15-5
Adding users to a group	15-6
Windows 200x Active Directory	15-7
Disabling Active Directory	15-7
Configuring the Storage Mirroring service for Active Directory	15-8
Appendix A: Storage Mirroring Commands	A-1
Conventions	A-1
Storage Mirroring Commands	A-2
Compression Disable	A-2
DTCL Scripting	A-45
Variables	A-45
Flow control	A-46
DTCL Return Values	A-47
DTCL Scripting Examples	A-51
Using batch files	A-51
Basic getting started script	A-51
Basic getting started script with failover	A-51
Basic restore script	A-52
Basic getting started script with variables	A-52
Basic flow control script	A-52
Appendix B: Application Failover	B-1
Script Sample	B-2
Appendix C: Disaster Recovery for MSCS	C-1
Cluster Configurations	C-2
MSCS and Storage Mirroring	C-3
Providing Disaster Recovery Using Storage Mirroring	C-4
Cluster to Cluster Configuration	C-5
Configuring the first cluster	C-5
Configure the virtual server on the first cluster	C-5
Install and Configure Storage Mirroring on the first cluster	C-9
Configure the Storage Mirroring replication set	C-9
Configuring the second cluster	C-10
Configure the virtual server on the second cluster	C-10
Install Storage Mirroring on the second cluster	C-10
Completing the first cluster configuration	C-10
Configure the Double-Take Source Connection resource	C-11
Bring the Components Online	C-12

Cluster to Standalone Configuration	C-13
Configuring the cluster	C-13
Configure the virtual server on the cluster	C-13
Install and Configure Storage Mirroring on the cluster	C-17
Configure the Storage Mirroring replication set	C-17
Configure the Double-Take Source Connection resource	C-18
Configuring the standalone target	C-20
Completing the cluster configuration	C-20
Bring the Components Online	C-20
Standalone to Cluster Configuration	C-21
Configuring the standalone sever	C-21
Configuring the cluster	C-21
Configure the virtual server on the cluster	C-21
Install Storage Mirroring on the cluster	C-25
Completing the standalone configuration	C-26
Understanding the Double-Take Source Connection resource	C-27
Appendix D: NAT Environments	D-1
NAT Overview	D-1
Storage Mirroring and NAT Overview	D-2
Modifying the Storage Mirroring Ports	D-4
Configuring the Router	D-6
Using the Management Console in a NAT Environment	D-7
Inserting the servers in the Management Console	D-7
Establishing the connection	D-8
Using the Text Client in a NAT Environment	D-9
Appendix E: Firewall Environments	E-1
Verifying Storage Mirroring Port Settings	E-1
Opening the Firewall Ports	E-2
Appendix F: Glossary	F-1
Index	I

Preface

Intended Audience

The intended audience for the HP OpenView Storage Mirroring® *User's Guide* is network administrators with expertise in Microsoft® Windows®. The network administrator should be responsible for setup and maintenance of the network and should have working experience installing, configuring, and maintaining network applications.

Organization of This Guide

Each chapter in this guide is organized in primarily the same manner. Each chapter, if applicable, has a topic overview and is then divided by each specific Storage Mirroring client. The overview contains information on the topic's features, functions, and options, while the remainder of the chapter explains how the topic works with respect to each client. For example, Chapter 7, *Mirroring*, is divided into *Mirroring Overview*, *Managing Mirroring Using the Management Console*, and *Managing Mirroring Using the Text Client*.

NOTE: For detailed installation instructions and quick setup instructions, see the Storage Mirroring *Getting Started* guide.

Conventions Used in This Guide

The following conventions are used throughout this guide:

- ◆ **Bold** is used for items you click or select with the mouse, such as menu names and items, dialog box options, or button names. For example, when you see **File**, **Options**, you should select the **File** menu and then choose **Options** from the pull-down menu.
- ◆ *Italics* is used for cross-references to other sections or chapters in this guide or to other books.
- ◆ Code text is used to indicate two types of items:
 - ◆ Text that is displayed onscreen exactly as shown, such as screen messages or error messages
 - ◆ Text that should be entered exactly as shown, such as directories, file names, or URLs
- ◆ **bold code text** is used for Double-Take Command Language (DTCL) command names. The following conventions are used to indicate command syntax for DTCL:
 - ◆ UPPERCASE letters are used for commands that must be typed as shown.
 - ◆ Underlined letters can be used as a shortcut for the command.
 - ◆ *lowercase italicized* letters are variables such as file names, user names, or machine names. These items can also be substituted with defined variables as described in *DTCL Scripting* on page A-45.
 - ◆ Angle brackets, < and >, surround required items that must be supplied with the command.
 - ◆ Square brackets, [and], surround optional items that can be supplied with the command but are not required.
 - ◆ The pipe character, |, separates items in a list
 - ◆ Identifiers that contain a space or non-alphanumeric characters must be enclosed in quotation marks. For example:
 - ◆ "domain.com"
 - ◆ "machine name"
 - ◆ "129.65.35.45"

Online References

The online manual and online help provide a means of accessing product information. The online help for the Management Console, Failover Control Center, and Text Client are automatically installed with the Storage Mirroring client components. The manuals are also automatically installed.

Online help

The online help requires a version 4 or later browser (version 4.5 or later is better). The online help can be accessed by:

- ◆ **Management Console and Failover Control Center**
 - ◆ Pressing the **F1** key
 - ◆ Clicking the **Help** button on screens where it is available
 - ◆ Selecting **Help, Help Topics**
- ◆ **Text Clients**
 - ◆ Typing `help`

Online manual

The online manuals require Adobe® Acrobat® Reader®. You can install Adobe Acrobat Reader by downloading the latest version from the Adobe web site at www.adobe.com. The online manuals can be accessed by opening `\docs\instlwin.pdf` or `\docs\dtuser.pdf`.

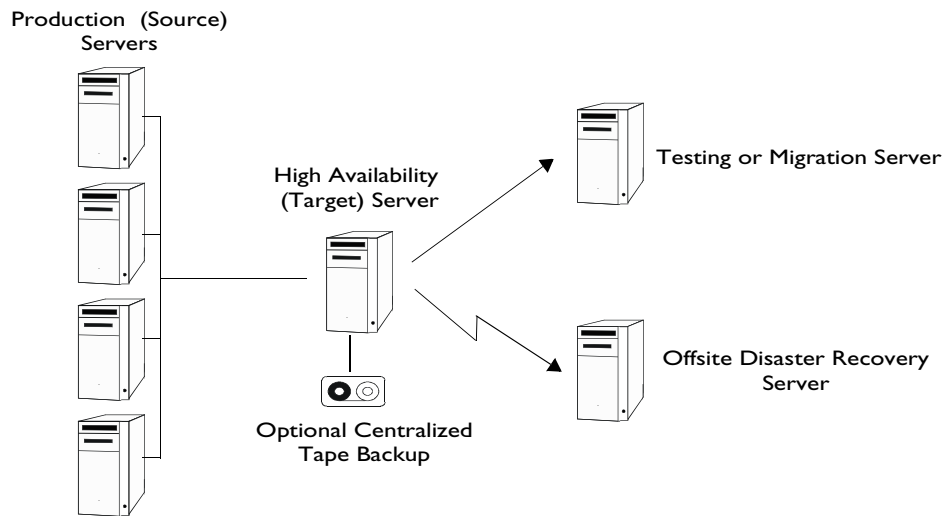
I ► Introduction to Storage Mirroring

Storage Mirroring is a real-time data replication and failover software product. Storage Mirroring augments your existing data protection strategy by reducing downtime and data loss, and it provides these services with minimal impact on existing network and communication resources.

Storage Mirroring allows you to specify mission-critical data that must be protected and replicates, in real-time, that data from a production machine, known as the source, to a backup machine, known as the target. The target machine, on a local network or at a remote site, stores the copy of the critical data from the source. Storage Mirroring monitors any changes to the critical data and sends the changes to the target machine. By replicating only the file changes rather than copying an entire file, Storage Mirroring allows you to more efficiently use resources.

Storage Mirroring lets you implement various data protection solutions including:

- ◆ Local high availability services
- ◆ Offsite disaster recovery services
- ◆ Enhanced centralized backup using third-party backup systems



Sample Configuration

Storage Mirroring Features

Storage Mirroring has high availability and disaster recovery management features, resource utilization features, multi-platform features, improved backup benefits, among other features.

High availability and disaster recovery management features

- ◆ **Continuous, transaction level replication** reduces your network's exposure to data loss by immediately replicating file changes to the target machine.
- ◆ **Failover capability** reduces downtime to seconds or minutes and provides automatic recovery of key business functions and data.
- ◆ **Flexible configuration options** include many-to-one, one-to-many, chained, and same server configurations across any IP based LAN or WAN.
- ◆ **Many-to-one capability** reduces the total cost of implementing Storage Mirroring for disaster recovery by allowing one target machine to support multiple source machines.
- ◆ **Extensive reporting and verification options** prevents silent failures by confirming that your data protection systems are working properly.

Resource utilization features

- ◆ **Selective file replication** allows you to select the critical files and directories that should be replicated. Compared to whole volume systems, this can dramatically reduce the bandwidth required for replication.
- ◆ **File update technology** transmits only the updates made to files, not whole disk blocks or whole files. This dramatically reduces the amount of bandwidth required and as a result, Storage Mirroring runs efficiently, even over low-bandwidth links.
- ◆ **Transmission control option** includes real-time or restricted transmission based on bandwidth usage (thus preserving bandwidth for other applications) or scheduled events (allowing you to specify criteria for Storage Mirroring transmission when real-time transmission is not practical). You can also restrict data flow to private links or a public network.
- ◆ **Loosely coupled communications** allows Storage Mirroring to operate asynchronously. Source machines continue processing even when the network link is congested.
- ◆ **Throughput Diagnostics Utility (TDU)** allows you to run Storage Mirroring in a monitor-only mode to track the amount of bandwidth that a connection will require.

Improved backup benefits

- ◆ **Open file mirroring** allows access to open files on the source machine while Storage Mirroring keeps a copy of data on the target machine. This capability eliminates the need to schedule source machine downtime for backing up open files.
- ◆ **Simplify backup management** by replicating data from multiple source machines to a single target machine. By backing up the target machine to a large backup device, you avoid the complexity of managing many smaller backup systems on the source machines.
- ◆ **Reduce production machine and network load** by backing up the target machine. Third-party backup software does not need to run or transfer data from the source machine when it is running on the target.

NOTE: Storage Mirroring does not replace traditional tape backup software. Tape backups are still recommended to preserve historical versions of files.

Other features and benefits

- ◆ **Simplified application testing** allows the replicated copy of data residing on the target machine to be a convenient source of data for software test operations. By copying selected sets of data to a test machine, network administrators can easily test application upgrades.
- ◆ **Operating system migration and testing** allows Storage Mirroring to be used to replicate data to a new machine during an operating system upgrade. Users can continue to operate on the current machine while applications, security, and connectivity are tested on the new system.
- ◆ **Data distribution** in a one-to-many configuration allows automatic distribution of changes from a master machine (the source) to multiple target machines on a LAN or WAN.

Storage Mirroring Operations

Storage Mirroring performs four basic types of operations:

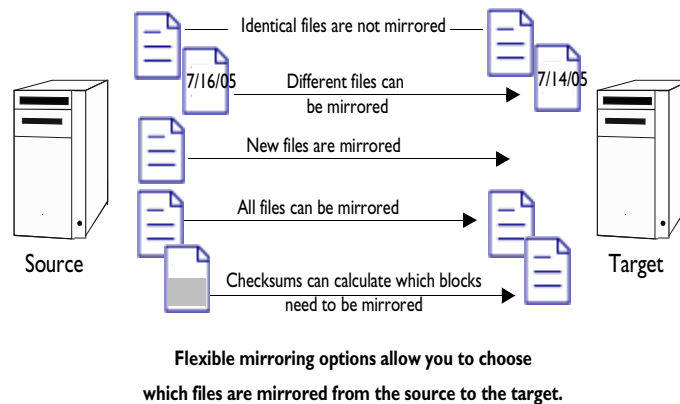
- ◆ **Mirroring**—The initial copy or subsequent resynchronization of selected data
- ◆ **Replication**—The on-going capture of byte-level file changes
- ◆ **Failure monitoring and failover**—The ability to monitor and stand-in for a machine, in the event of a failure
- ◆ **Restoration**—A mirror of selected data from the target back to the source

Each operation is briefly described in the following sections. For complete details, see the corresponding chapter.

Mirroring

Mirroring is the process of transmitting user-specified data from the source to the target so that an identical copy of data exists on the target. When Storage Mirroring initially performs mirroring, it copies all of the selected data including file attributes and permissions. Mirroring creates a foundation upon which Storage Mirroring can efficiently update the target machine by replicating only file changes.

If subsequent mirroring operations are necessary, Storage Mirroring can mirror specific files or blocks of changed data within files. By mirroring only files that have changed, network administrators can expedite the mirroring of data on the source and target machines.




Mirroring has a defined end point - when all of the selected files from the source have been transmitted to the target. When a mirror is complete, the target contains a copy of the source files at that point in time.

Replication

Replication is the real-time transmission of file changes. Unlike other related technologies, which are based on a disk driver or a specific application, the Storage Mirroring replication process operates at the file system level and is able to track file changes independently from the file's related application. In terms of network resources and time, replicating changes is a more efficient method of maintaining a real-time copy of data than copying an entire file that has changed.

After a source and target have been connected through Storage Mirroring, file system changes from the user-defined mission-critical data are tracked. Storage Mirroring immediately transmits these file changes to the target machine. This real-time replication keeps the data on the target machine up-to-date with the source machine and provides high availability and disaster recovery with minimal data loss.



Unlike mirroring which is complete when all of the files have been transmitted to the target, replication captures the changes as they are written to the source and sends them to the target continuously. Replication keeps the target up-to-date and synchronized with the source.



Failure monitoring/failover

Failover is a process in which a target machine stands in for a failed source machine. As a result, user and application requests that are directed to the failed source machine are routed to the target machine.

Storage Mirroring monitors the status of machines by tracking network requests and responses exchanged between source and target machines. When a monitored machine misses a user-defined number of requests, Storage Mirroring assumes that the machine has failed. Storage Mirroring then prompts the network administrator to initiate failover, or, if configured, it occurs automatically.

The failover target assumes the network identity of the failed source. When the target assumes the identity of the source, user and application requests destined for the source machine or its IP address(es) are routed to the target. After failover, user and application requests directed to the source machine are routed to the target machine.

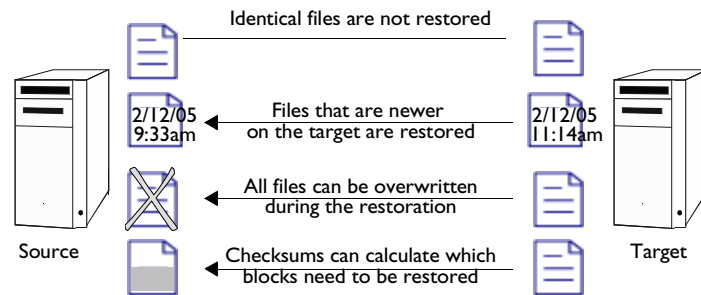
When partnered with the Storage Mirroring data replication capabilities, failover routes user and application requests with minimal disruption and little or no data loss. In some cases, failover may be used without data replication to ensure high availability on a machine that only provides processing services, such as a web server.



Restoration

Restoration provides an easy method for copying replicated data from the target back to its original location on the source. The process only requires you to select the source, target, and the appropriate replication set. There is no need to select files or to remember where the data came from on the source since that information is maintained by Storage Mirroring.

Restoration can be used if the source data is lost due to a disk crash or when the most up-to-date data exists on the target due to failover. At the time of a source machine failure, your Storage Mirroring target will contain the same data as your Storage Mirroring source. If you are using the Storage Mirroring failover capabilities, users can continue updating data on the target machine while the problems on the source are resolved. Because of the continued updates on the target, when the source machine is ready to come back online, the two machines will no longer contain the same data. Restoration is the process of copying the up-to-date data from the target back to the original source or new source location when bringing the source back online.



Flexible restoration options allow you to choose which files are restored from the target to the source.

Like mirroring, restoration has a defined end point - when all of the selected files from the target have been transmitted back to the source. When a restoration is complete, the source and target are again synchronized.

Sample Configurations


Storage Mirroring is an exceptionally flexible product that can be used in a wide variety of network configurations. To implement Storage Mirroring effectively, it is important to understand the possible configuration options and their relative benefits. Storage Mirroring configuration options can be used independently or in varying combinations.

NOTE:

The Storage Mirroring replication and failover features have different machine requirements. If you use Storage Mirroring to monitor and stand in for failed machines, HP recommends that the source and target machines have identical versions of operating system and applications. By contrast, Storage Mirroring replication operations do not require that the source and target machines have identical versions of operating system and applications.

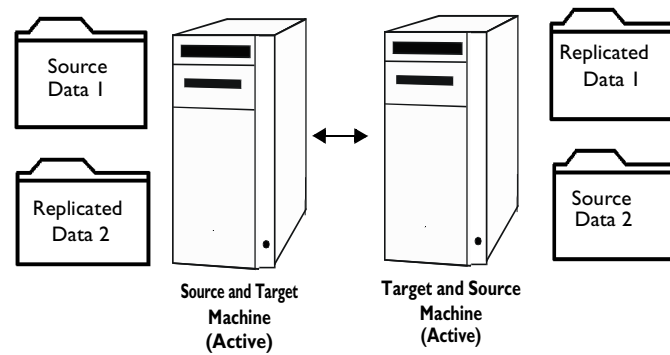
If you are using Microsoft Cluster Server, a dual-node cluster can act as any machine for any of the following configuration samples. For more information on using Storage Mirroring with Microsoft Cluster Server, see [Disaster Recovery for MSCS](#) on page C-1.

One-to-One, Active/Standby




Description	One target machine, having no production activity, is dedicated to support one source machine. The source is the only machine actively replicating data.
Applications	This configuration is appropriate for offsite disaster recovery, failover, and critical data backup. This is especially appropriate for critical application servers such as Exchange, SQL Server, and web servers. This is the easiest configuration to implement, support, and maintain.
Considerations	This configuration requires the highest hardware cost because a target machine is required for every source machine.

One-to-One, Active/Active




Description	Each machine acts as both a source and target actively replicating data to each other.
Applications	This configuration is appropriate for failover and critical data backup. This configuration is more cost-effective than the Active/Standby configuration because there is no need to buy a dedicated target machine for each source. In this case, both machines can do full-time production work.
Considerations	<p>Coordination of the configuration of Storage Mirroring and other applications can be more complex than the Active/Standby configuration.</p> <ul style="list-style-type: none">◆ During replication, each machine must continue to process its normal workload.◆ Administrators must avoid selecting a target destination path that is included in the source's replication set. Any overlap will cause an infinite loop.◆ To support the production activities of both machines during failover without reducing performance, the failover target machine should have sufficient disk space and processing resources.◆ Failover and failback scripts must be implemented to avoid conflict with the existing production applications.

Many-to-One




Description	Many source machines are protected by one target machine.
Application	This configuration is appropriate for offsite disaster recovery. This is also an excellent choice for providing centralized tape backup because it spreads the cost of one target machine among many source machines.
Considerations	The target machine must be carefully managed. It must have enough disk space and RAM to support replication from all of the source systems. If using failover, scripts must be coordinated to ensure that, in the event that the target machine stands in for multiple machines, the respective applications will not conflict. The NIC(s) must be able to accommodate traffic for all monitored machines if they should fail simultaneously.

One-to-Many



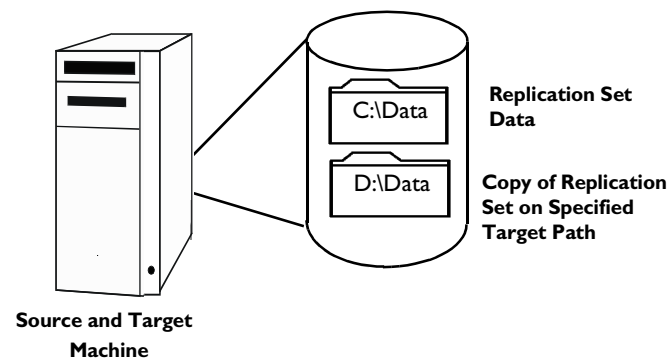
Description	One source machine sends data to multiple target machines. The target machines may or may not be accessible by one another.
Applications	This configuration provides offsite disaster recovery, redundant backups, and data distribution. For example, this configuration can replicate all data to a local target machine and separately replicate a subset of the mission-critical data to an offsite disaster recovery machine.
Considerations	Updates are transmitted multiple times across the network. If one of the target machines is on a WAN, the source machine is burdened with WAN communications.

Chained



Description	One or more source machines sends replicated data to a target machine that in turn acts as a source machine and sends selected data to a final target machine which is often offsite.
Applications	This is a convenient approach for integrating local high availability with offsite disaster recovery. This configuration moves the processing burden of WAN communications from the source machine to the target/source machine.
Considerations	The target/source machine could become a single point of failure for offsite data protection.

Single Machine



Description	Source and target Storage Mirroring components are loaded on the same machine allowing data to be replicated from one location to another on the same volume or to a separate volume on the same machine. These could be locally attached SCSI drives or Fibre Channel based SAN devices.
Applications	This configuration is useful for relocating data within existing machines so that open files can be backed up. If the source and target copies of the data are located on different drives, this configuration does support high availability of the replication set in the event that the source hard drive fails.
Considerations	This configuration does not provide high availability for the entire machine. This configuration must be configured carefully so that an infinite loop is not created.

2 ► Management Console

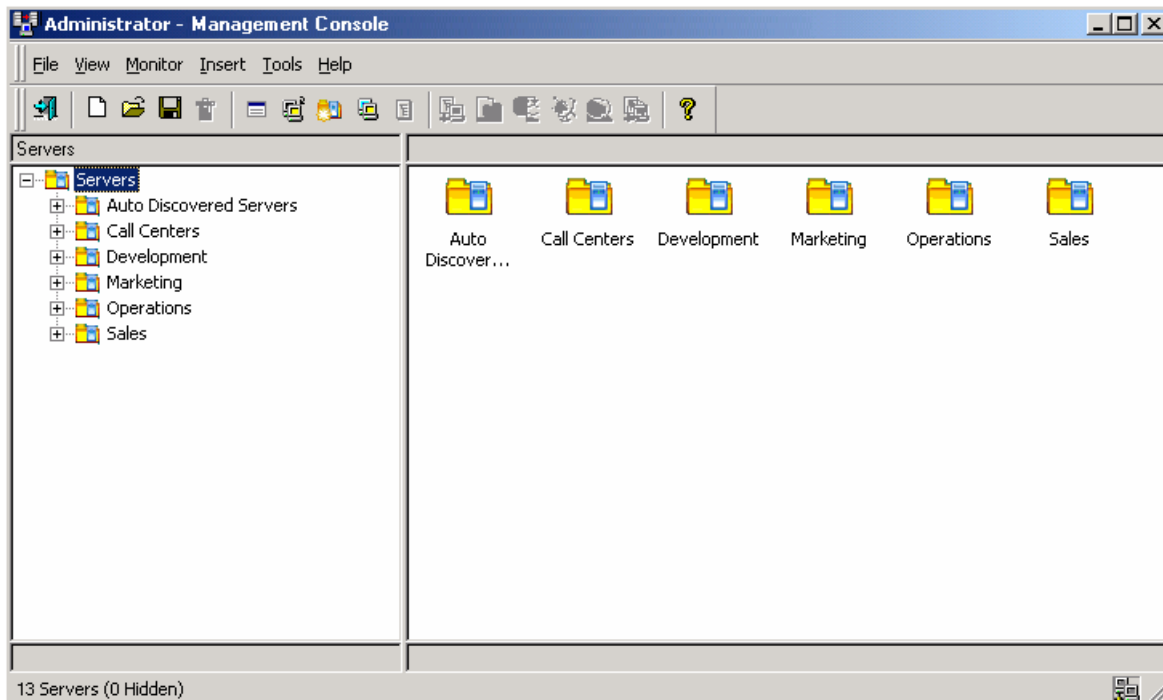
The Management Console is a 32-bit GUI client that can manage and monitor Storage Mirroring components on any platform. Storage Mirroring machines are displayed along with active connection information. The Management Console can also configure failover settings.

NOTE: You may not have access to some of the components or see certain display options if you are using a newer version of the Management Console to control an older version of your source or target.

Starting the Management Console

From the machine that you installed the Storage Mirroring client, select **Start, Programs, Storage Mirroring, Management Console**.

The Management Console is divided into a left and right pane. The views in the panes change depending on what is highlighted. For example, in the left pane when the root of the tree, labeled Storage Mirroring Servers, is highlighted, all of the machines running Storage Mirroring are displayed in the right pane. If you double-click on the root, those same servers expand or collapse in the branches of the tree. More detailed information can be found in [Management Console Tools](#) on page 13-2.














Understanding the Display

The left pane is a tree view displaying each machine running Storage Mirroring. As each machine name is highlighted, note the following items:




- ◆ The status bar at the top of the pane displays the modules loaded on each machine: Source, Target, Source and Target, or None. (None indicates that the core functionality of Storage Mirroring is running, but neither the source nor target modules have been loaded.)
- ◆ The icons in the left pane indicate what modules are loaded on each machine.

Left Pane Icons

Icon	Description
	A blue server in front of a yellow folder indicates a group.
	The group icon with a yellow triangle and a black exclamation point indicates there is a communication error with one of the servers in the group.
	The group icon with a red circle with a white X indicates there is a connection error with one of servers in the group.
	A blue machine indicates a Storage Mirroring source.
	A yellow machine indicates a Storage Mirroring target.
	Yellow and blue machines indicate a Storage Mirroring source and target.
	A single machine with a hammer indicates that neither the source nor target modules are loaded.
	A machine with two red vertical lines over it indicates that the target module is paused.
	A machine with red horizontal lines over it indicates that a restore may be required because the target machine is or has been standing in for the source machine due to failover.
	Any of the machine icons can appear with a red X. The red X indicates: <ul style="list-style-type: none">◆ The Management Console cannot communicate with that Storage Mirroring machine.◆ There is a problem with an established Storage Mirroring connection. Look in the right pane of the Management Console or log on to the machine in the left pane to determine the exact problem.
	A machine icon can appear with a black X. The black X indicates that the machine is not running Storage Mirroring.




- ◆ The icon to the left of the machine name indicates the security access that is granted for that machine. For detailed information on security and access rights, see [Security](#) on page 15-1.

Security Icons

Icon	Description	Access Granted
	This icon is a computer with a gear and it indicates the Storage Mirroring security is set to administrator access.	Administrator rights
	This icon is a computer with a magnifying glass and it indicates the Storage Mirroring security is set to monitor only access.	Monitor rights
	This icon is a lock and it indicates the Storage Mirroring security is set to no access.	No rights

- ◆ After you have logged into a machine, a collapsible and expandable tree view of the volumes and directories it contains is displayed. To expand the tree, click on the plus sign icon. To collapse the tree, click on the minus sign icon.

- ◆ Under each machine, you can create replication sets. Different icons represent different replication set states.

Icon	Description
	A red tree icon indicates a modified replication set that has not yet been saved.
	A black tree icon indicates a saved replication set.
	A blue tree icon indicates that the replication set has been modified in two separate instances of the Management Console and the modifications do not match.

- ◆ The right pane displays the following information depending on the item selected in the left pane of the Management Console and the tab selected on the right pane.


















Right Pane Display

Left Pane	Right Pane	Description
Group	No Tabs	All groups and servers contained in the selected group are displayed.
Source Machine	Source Tab	The active connections, if any, from that source machine are displayed.
Source Machine	Target Tab	The display is blank.
Target Machine	Source Tab	The display is blank.
Target Machine	Target Tab	The active connections, if any, to that target machine are displayed. ^a The server name is also displayed, if it is in the Management Console tree. If it is not in the tree, the IP address is displayed.
Replication Set	No Tabs	The available target machines are displayed so that you can drag and drop a replication set onto a target to establish a connection.
Volume or Directory	No Tabs	The directories and files contained in that volume or directory are displayed.

a.If transmission is stopped (manually stopped, outside of a scheduled transmission window, network error between the source and target, and so on) the replication set will not be displayed on the Target tab. When transmission (re)starts, the replication set will (re)appear.

Management Console toolbar


The following icons appear in the Management Console toolbar.

Icon	Function
	Exit the Storage Mirroring Management Console
	Create a new workspace
	Open an existing workspace
	Save the active workspace
	Remove the selected item (group, server, or replication set)
	Add a message window
	Launch Failover Control Center
	Insert a group
	Insert a server
	Insert a replication set
	Open the Connection Manager
	Start the mirroring process
	Start the replication process
	Start the transmission process
	Start the verification process
	Disconnect the current connection
	Display Storage Mirroring program information

Logging On and Off of Storage Mirroring

To ensure protection of your data, Storage Mirroring offers multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine running Storage Mirroring. To gain access to a particular Storage Mirroring source or target, the user must provide a valid operating system user name and password and the specified user name must be a member of one of the Storage Mirroring security groups. Once a valid user name and password has been provided and the Storage Mirroring source or target has verified membership in one of the Storage Mirroring security groups, the user is granted appropriate access to the source or target and the corresponding features are enabled in the client. Access to Storage Mirroring is granted on one of the following three levels:

- ◆ **Administrator Access**—All Storage Mirroring features are available for that machine. For example, this access level includes creating replication sets and establishing Storage Mirroring connections.
 - ◆ **Monitor Access**—Statistics can be viewed on that machine, but Storage Mirroring features are not available. For example, this access level does not allow the user to create or modify replication sets or create or modify Storage Mirroring connections, but does allow you to view the connection statistics for any established Storage Mirroring connections on that machine.
 - ◆ **No Access**—The machine appears in the Storage Mirroring Management Console and can be pinged from the Storage Mirroring Text Client, but no other access is available.
1. Highlight a machine on the left pane of the Management Console. By double-clicking the machine name, Storage Mirroring automatically attempts to log you on to the selected machine. Verify your access by the resulting icon.
 2. If you have no access, the Logon dialog box will automatically appear. If you have monitor access or want to log on with a different username, right-click the machine name and select **Logon**.



3. Specify your **Username**, **Password**, **Domain**, and whether you want your password saved. Click **OK**.

NOTE: If your activation code is missing or invalid, you will be prompted to open the Server Properties general tab to add or correct the code. Select **Yes** to open the Server Properties dialog box or select **No** to continue without adding an activation code.

4. Verify your access by the resulting icon and log on again if necessary.
5. To log off of a Storage Mirroring machine, right-click the machine name on the left pane of the Management Console and select **Logout**.

NOTE: For detailed information on Storage Mirroring security, see [Security](#) on page 15-1.

If the login does not complete within 30 seconds, it is automatically canceled. If this timeout is not long enough for your environment, you can increase it by adjusting the **Communication Timeout** on the Configuration tab of the Management Console properties. Select **File, Options**, from the Management Console to access this screen.

Storage Mirroring uses ICMP pings to verify server availability during the login process. If your Storage Mirroring server is across a router or firewall that has ICMP pings disabled, you will need to disable the Storage Mirroring ICMP ping verification. To do this, select **File, Options**, from the Management Console and disable **Use ICMP to verify server availability**.

Stopping the Service

To stop the Storage Mirroring service from the Management Console, right-click the machine name on the left pane of the Management Console and select **Shutdown**.

Storage Mirroring Ports


Both the Management Console and the Storage Mirroring server use ports to communicate with each other.

Management Console ports

The Storage Mirroring Management Console uses two ports for Storage Mirroring communications.

- ◆ **Heartbeat Port**—The Management Console uses this port to listen for heartbeats transmitted from Storage Mirroring servers.
- ◆ **Status Transmit Port**—The Management Console uses this port to send requests to Storage Mirroring servers so that it can update its at-a-glance monitoring information.

1. To access the port settings, select **File, Options**.



2. Under Network on the **Configuration** tab, modify the **Status Transmit Port** and specify the number that you want to assign to the port. Under Automatic Service Discovery on the **Configuration** tab, modify the **Heartbeat Port** and specify the number that you want to assign to the port. For both ports, the default values are 1105 and 1100, respectively. The valid values are 1025 to 65535.
3. Click **OK** to save the settings.

NOTE: These changes will add any new Storage Mirroring machines to the display, but will not remove Storage Mirroring machines that are using the previous port settings. To remove those machines, the Management Console must be restarted.

Storage Mirroring server ports

The Storage Mirroring service uses three ports for Storage Mirroring communications.

- ◆ **Service Listen Port**—The Storage Mirroring server uses this port to listen for commands from other Storage Mirroring servers.
- ◆ **Heartbeat Transmit Port**—The Storage Mirroring server uses this port to transmit heartbeats.
- ◆ **Status Listen Port**—The Storage Mirroring server uses this port to listen for requests to update at-a-glance monitoring from the Storage Mirroring Management Console.

1. Right-click a machine on the left pane of the Management Console and select **Properties**.

The screenshot shows the 'Server Properties' dialog box with the 'Network' tab selected. The dialog has a title bar with a close button. Below the title bar are tabs for 'Target', 'Database', 'Logging', 'E-mail Notification', 'General', 'Setup', 'Network', 'Queue', and 'Source'. The 'Network' tab is active. A warning icon and text state: 'Changing the values on this page will not take effect until the Storage Mirroring service has been restarted on the specified system.' The 'Interface' section contains: 'Default Address' (None), 'Default Protocol' (TCP/IP), 'Service Listen Port' (1100), 'Heartbeat Transmit Port' (1100), and 'Status Listen Port' (1105). The 'Heartbeat' section contains: 'Time between Heartbeats (seconds)' (3) and 'Missed Heartbeat Limit' (10). At the bottom are 'OK', 'Cancel', and 'Help' buttons.

2. On the **Network** tab, modify the **Service Listen Port**, **Heartbeat Transmit Port** and/or the **Status Listen Port** and specify the number that you want to assign to the ports. The default values are 1100, 1100 and 1105, respectively. The valid values are 1025 to 65535.
3. Click **OK** to save the settings.

NOTE: You must stop and restart the Storage Mirroring service for this change to take effect.

Managing the Management Console Tree

To better manage the servers that appear in the Management Console, you can customize the server display to fit your needs. You can create groups and move servers to those groups to help you organize your environment. Within the groups, you can insert, remove, hide or unhide servers. Each of these functions is detailed in the following sections.

Groups

The left pane of the Management Console is a tree view of the Storage Mirroring servers. By default, the first group in the tree is the Auto-Discovered Servers group. All Storage Mirroring servers that are automatically discovered will be added to this group. Use server groups in a hierarchical structure to help you organize your environment.

Creating groups

Use one of the following methods to create a new group:

- ◆ Right-click anywhere on the left pane of the Management Console and select **New, Group**.
- ◆ Right-click on a group icon on the right pane of the Management Console and select **New, Group**.
- ◆ Click Add Group on the toolbar.
- ◆ Use the menu bar and select **Insert, Group**.

The location of the new group that is created will depend on what was highlighted. If the root of the tree was highlighted, the new group will be created as a child of the root. If a group or server within a group was highlighted, the new group will be created as a child of that group. Name the newly inserted group with a unique name by typing over the default name and pressing Enter. This process is similar to naming a new folder in Windows Explorer. You can also rename an established group by double-clicking on the existing name. Type the new group name over the existing name and press Enter.

Removing groups

Use one of the following methods to remove a group.

- ◆ Right-click on a group on the left of the Management Console and select **Remove**.
- ◆ Right-click on a group on the right pane of the Management Console and select **Remove**.
- ◆ Click Remove Item on the toolbar.
- ◆ Highlight a group and press the Delete key.

You will be prompted to confirm the removal of the group and its subgroups and any servers (displayed or hidden) contained in them. Click **OK**.

If auto-discovery is enabled on the Management Console, those servers that also have auto-discovery enabled will automatically be repopulated back in the default Auto-Discovered Servers group. If auto-discovery is disabled on the Management Console or for individual servers, servers will need to be manually inserted into the Management Console.

NOTE: The Remove toolbar button also removes servers and replication sets, so make sure you have the correct item highlighted before clicking the toolbar button.

You cannot remove the default Auto-Discovered Servers group.

Servers

Within your server groups, you do have the ability to further manage the servers that are displayed by using the following functions.

- ◆ **Moving Servers**—Servers that are auto-populated can be moved to different groups within the Management Console tree.
- ◆ **Insert**—If a machine is not displayed on the Management Console, it can be manually inserted. This feature is useful for machines that are across a router or on a different network segment. These types of machines will not automatically appear in the Management Console because the UDP heartbeat is not passed by the router.

NOTE: If a machine is manually inserted into the Management Console, it will automatically be saved in your workspace and will appear the next time the Management Console is started.

-
- ◆ **Remove**—If you do not want to see a server in the Management Console, it can be permanently removed from the display. You might need to remove a server that was manually added, if that server is no longer needed. Or if there are servers within the network that another administrator is responsible for, you can remove them from your display.

NOTE: If Storage Mirroring heartbeats are enabled or a server is listed in Active Directory, a removed server will automatically be added back to the server list the next time a heartbeat is received. You need to disable heartbeats to permanently remove it.

- ◆ **Hide**—If you do not want to see a server in the Management Console and do not want to disable heartbeats, you can hide the server from view. This keeps the server in the Management Console's internal list of servers, but does not display it in the server tree, any dialog boxes, or any field/menu selections.
- ◆ **Unhide**—You can unhide, or display, a hidden server at any time you want to access that server. The server will be displayed in the server tree, all dialog boxes, and field/menu selections.

Moving servers

You can move servers to groups by either of the following methods

- ◆ Drag and drop a server into the desired group. With this method you can move one server at a time in the left pane of the Management Console. You can move more than one server at a time by using this method from the right pane of the Management Console.
- ◆ Insert a server into the desired group. Inserting an existing server to the tree will move the first occurrence to that new location.

A Storage Mirroring server will only appear once within the entire Management Console tree. Servers cannot be placed into multiple groups.

Insert a server

Use the following instructions to insert a server into the Management Console.

NOTE: A server that already exists in the tree will be moved to the currently selected group if you attempt to insert it again.

1. Right-click on a group and select **New, Server** or highlight a group and select **Insert, Server**.
2. Type the machine name or IP address and the port number, if it is different than the default.
3. Select **Test** to determine if the machine is available and running Storage Mirroring. Depending on the results of the test, you will see one of the following:
 - ◆ If the machine was located and Storage Mirroring is running on it, the machine name will be inserted into the Management Console.
 - ◆ If the machine was located, but Storage Mirroring is not running on it, you will be prompted that the machine was found, but Storage Mirroring is not available. You will be given the choice of inserting the machine anyway. Click **Yes** to insert the server into the Management Console or click **No** to return to the Insert Server dialog box. Enter another server name or click **Cancel** to return to the Management Console without inserting a server.
 - ◆ If the server could not be located, you will see a message to the right of the **Test** button stating that it was not found. Enter another server name or click **Cancel** to return to the Management Console without inserting a server.

NOTE: At any time while Storage Mirroring is attempting to locate the machine, click **Stop** to cancel the test. If you do not manually test a machine before inserting it, Storage Mirroring will automatically test it for you.

Remove a server

To remove a server, right-click on the server in the left or right pane of the Management Console and select **Remove**. You can also select Remove Item from the toolbar.

If auto-discovery is enabled on the Management Console, those servers that also have auto-discovery enabled will automatically be repopulated back in the default Auto-Discovered Servers group. If auto-discovery is disabled on the Management Console or for individual servers, servers will need to be manually inserted into the Management Console.

Hide a server

To hide a server, right-click on a server in the left or right pane of the Management Console and select **Hide**.

NOTE:	<p>If you attempt to insert a server that is already in the tree but hidden, you will be prompted to unhide the server and insert it into the selected group.</p> <p>Be careful if you hide a server with an established Storage Mirroring connection. If that connection goes into an error state, you will not be able to see the connection in the Management Console. The Storage Mirroring log, Event Viewer, and other monitoring methods will still be functioning to alert you to the error. Hiding the server only removes it from the Management Console display.</p> <p>If a target server with an established Storage Mirroring connection is hidden and you open the Connection Manager for that connection via the source, you will see the target and IP address displayed in the Target and Route fields, respectively. This is the only time you will see a hidden server.</p>
--------------	---

Unhide a server

To unhide, or display, a hidden server, you can insert the server or use the following instructions.

1. Select **View, Unhide Servers**.
2. Select one or more servers by using Ctrl-click or Shift-click. You can also click **Select All** to select all of the servers in the list.
3. Click **Unhide**.

Sharing group and server configuration

All of the group and server information is stored on the local machine. When you close the Management Console, the group information is saved and will be persisted the next time you open the Management Console on this same machine. If you want to share the group configuration with another user, you can export the group configuration information (**File, Export server group configuration**) to an .xml file which is saved to the same directory where a Storage Mirroring is installed. That file can then be copied and imported (**File, Import server group configuration**) on another machine.

Storage Mirroring Workspaces

The Management Console workspace contains the display of the panes of the Management Console and any servers that may have been inserted. Multiple workspaces can be used to help organize your environment or to view settings from another machine.

Saving a workspace

As you size, add, or remove windows in the Management Console, you can save the workspace to use later or use on another Storage Mirroring client machine.

Select **File** and choose:

- ◆ **Save Workspace** to save the current workspace. If you have not previously saved this workspace, you must specify a name for this workspace.
- ◆ **Save Workspace As** to prompt for a new name when saving the current workspace.

Opening a workspace

From the Management Console, you can open a new workspace or open a previously saved workspace.

Select **File** and choose:

- ◆ **New Workspace** to open an untitled workspace with the default Storage Mirroring window settings.
- ◆ **Open Workspace** to open a previously saved workspace.

Getting Help for the Management Console

Context-sensitive help is available in the Management Console by:

- ◆ Clicking **Help** when it appears on various dialog boxes
- ◆ Pressing the **F1** key
- ◆ Selecting **Help**, **Help Topics**

Exiting the Management Console

To exit the Management Console, select **File**, **Exit** or select the **Exit** button on the toolbar.

3 ► Text Clients

The Double-Take Command Language (DTCL) is a scripting language that can be used in either the Text Client or Command Line Client to manage and monitor Storage Mirroring components. It can also be used in script files to execute series and combinations of commands to meet specific needs.

Because the Text Client and Command Line Client use the same DTCL commands, all examples use the Text Client. For a complete listing of the DTCL commands and the conventions used to document them, see [Storage Mirroring Commands](#) on page A-1.

The Text Client

The Text Client is a full-screen, text-based client that can be run from a command prompt. It can manage and monitor Storage Mirroring components on any platform regardless of the Text Client platform. The interface uses DTCL commands entered one line at a time. The Text Client offers command line editing capabilities, such as backspace cursor movement and forward/backward command history scrolling. All connection information is displayed in columnar text format on the screen.

Starting the Text Client

From the Windows desktop, select **Start, Programs, Storage Mirroring, Text Client** or from a command prompt, type the command `DTText`.

NOTE: The `DTText` command name is not case-sensitive.

Using the Text Client

The Text Client is divided into three sections:

- ◆ The top section displays source, connection, and statistical data. If specifying a command that requires output, like the `status` command, the output is also displayed in this section.
- ◆ The middle section is where the DTCL commands are entered.
- ◆ The lower section displays the DTCL command that was entered with any resulting messages. In this section, the security access granted is also displayed.

To use Storage Mirroring from the Text Client, enter the DTCL commands at the **Command** prompt in the middle of the Text Client screen.

```
Source Machine:           Storage Mirroring version 4.4

=====
Command:  login indy administrator ***** domain_name
=====
User access level set to DT_FULL_ACCESS
```

The Command Line Client

The Command Line Client can be run from a command prompt. It can manage and monitor Storage Mirroring components on any platform regardless of the Command Line Client platform. The Command Line Client is run from the operating system command line and the DTCL commands are run from the Storage Mirroring command line.

Starting and using the Command Line Client

There are three different methods of executing commands from the Command Line Client:

- ◆ **Interactive Entry**—At the directory prompt where Storage Mirroring is installed, type the command `DTCL -i`. A DTCL Command prompt will appear and the commands can be entered from that prompt. Any resulting errors are immediately displayed.

Interactive Entry

```
C:\Program Files\OpenView\Storage Mirroring> dtcl -i
Command: login indy administrator ***** domain_name
User access level set to DT_FULL_ACCESS
Command:
```

- ◆ **File Entry**—Create a file with all of the DTCL commands you want to run. At the directory prompt where Storage Mirroring is installed, type `DTCL -f filename` where `filename` is the name of the file containing the DTCL commands. For example, you might use the following `dtcl.txt` file to log on to a machine and display its replication sets.

File Entry

```
C:\Program Files\OpenView\Storage Mirroring> dtcl -f dtcl.txt
User access level set to DT_FULL_ACCESS
- List of rep sets -
Exchange                               enabled
C:\Program Files\OpenView\Storage Mirroring>
```

DTCL.TXT Used in File Entry

```
login indy administrator password domain_name;
source indy;
repset list;
```

- ◆ **Single Line Entry**—Determine all of the DTCL commands you want to run and enter them in a single line at the command prompt where Storage Mirroring is installed. For example, you might use the following command to log on to a machine and display its replication sets.

Single Line Entry

```
C:\Program Files\OpenView\Storage Mirroring> dtcl login indy administrator
password; source indy; repset list
User access level set to DT_FULL_ACCESS
- List of rep sets -
Exchange                               enabled
C:\Program Files\OpenView\Storage Mirroring>
```

NOTE: Because CMD.exe strips two layers of quotation marks during processing, any DTCL command that you use that requires quotation marks must have three quotation marks around it. For example, if your machine name had a space in it, `login "machine name" username password` would be sufficient for the Text Client or the Command Line Client interactive entry or file entry. But for the Command Line Client single line entry, you would have to use `login """"machine name"""" username password`.

Logging on and off of Storage Mirroring

To ensure protection of your data, Storage Mirroring uses native operating system security and requires a valid system username and password. If you do not have a valid system username and password for a selected machine, you will not be able to configure Storage Mirroring for that machine.

As you enter DTCL commands, Storage Mirroring automatically attempts to log on to the source or target machine. You can verify your access by the resulting message.

- ◆ **DT_Full_Access**—You have administrator rights to the selected Storage Mirroring machine.
- ◆ **DT_Monitor_Only_Access**—You have monitor rights to the selected Storage Mirroring machine.
- ◆ **DT_No_Access**—You do not have any rights to the selected Storage Mirroring machine.

You can also verify your access at the top of the Text Client screen depending on the commands you have entered and the information displayed in the upper half of the screen.

If you have monitor access or no access and want to log on as the administrator, use the `login` command. Log off of a machine by using the `logout` command.

Command	LOGIN
Description	Log on to a Storage Mirroring machine
Syntax	LOGIN <machine> <username> <password> [domain]
Options	<ul style="list-style-type: none">◆ machine—Name of the machine◆ username—Name of the user◆ password—Password associated with username.◆ domain—If logging in using a domain account, this is the domain name. If logging in using a local account, this is the machine name.
Examples	login indy administrator *****
Notes	<ul style="list-style-type: none">◆ The <code>login</code> command is not available when scrolling through the Text Client command history.◆ If characters in the password include non-alphanumeric characters, the password field must be enclosed in quotation marks.◆ The password cannot be a Storage Mirroring keyword. These are any DTCL command (source, target, and so on.) or any DTCL shortcut command (env, mon, rep, and so on).

Command	LOGOUT
Description	Logs off of a Storage Mirroring machine
Syntax	LOGOUT <machine>
Options	machine —Name of the machine
Examples	logout indy

Stopping the Service

To stop the Storage Mirroring service from the Text Client, use the `shutdown` command.

Command	<code>SHUTDOWN</code>
Description	Stops the Storage Mirroring service
Syntax	<code>SHUTDOWN <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>shutdown indy</code>

Specifying Source and Target Functionality

If you did not specify your machine as a source or target when starting the Storage Mirroring you will need to load the modules using the DTCL commands `load source` or `load target`.

Command	<code>LOAD SOURCE</code>
Description	Loads the Storage Mirroring source module
Syntax	<code>LOAD SOURCE <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>load source indy</code>

Command	<code>LOAD TARGET</code>
Description	Loads the Storage Mirroring target module
Syntax	<code>LOAD TARGET <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>load target jersey</code>

To remove a specific functionality, use the `unload source` or `unload target` commands.

Command	<code>UNLOAD SOURCE</code>
Description	Unloads the Storage Mirroring source module
Syntax	<code>UNLOAD SOURCE <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>unload source indy</code>

Command	<code>UNLOAD TARGET</code>
Description	Unloads the Storage Mirroring target module
Syntax	<code>UNLOAD TARGET <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>unload target jersey</code>

Storage Mirroring Ports

Both the text clients and the Storage Mirroring server use ports to communicate with each other.

Text client ports

The Storage Mirroring text clients use two ports for Storage Mirroring communications.

- ◆ **netport**—The text clients use this port to send commands to a Storage Mirroring server.
 - ◆ **unetport**—The text clients use this port to listen for heartbeats sent from a Storage Mirroring server.
1. To determine the current netport setting on a Storage Mirroring text client, use the `get` command with the `netport` option. The default port number is 1100.

Command	<code>GET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	<code>GET <option> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">◆ option—See Storage Mirroring Text Client program settings on page 14-15.◆ machine—Name of the machine
Examples	<ul style="list-style-type: none">◆ <code>get netport</code>◆ <code>get VerifyLogName</code>◆ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. To modify the netport setting on a Storage Mirroring text client, use the `set` command with the `netport` option. The `netport` option can be any number from 1025-65535 that you want to assign to the port.

Command	<code>SET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	<code>SET <option>=<value> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">◆ option—See Storage Mirroring Text Client program settings on page 14-15.◆ value—See Storage Mirroring Text Client program settings on page 14-15.◆ machine—Name of the machine
Examples	<ul style="list-style-type: none">◆ <code>set netport=1100</code>◆ <code>set VerifyLogName="HA server.log"</code>◆ <code>set DefaultAddress="129.12.34.3"</code>
Notes	<ul style="list-style-type: none">◆ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed.◆ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

```
Source Machine: indy          Storage Mirroring version 4.4          (Access Level: FULL)

=====
Command:
=====
> setlocal netport = 1110
```

3. To determine the current unetport setting, use the `get` command with the `unetport` option. The default port number is 1100.
4. To modify the unetport, use the `set` command with the `unetport` option. The `unetport` option can be any number from 1025-65535 that you want to assign to the port.

Storage Mirroring server ports

The Storage Mirroring service uses two ports for Storage Mirroring communications.

- ◆ **netport**—The Storage Mirroring server uses this port to listen for commands from a Storage Mirroring text client.
- ◆ **unetport**—The Storage Mirroring server uses this port to send heartbeats to Storage Mirroring clients.

1. To determine the current netport setting on a Storage Mirroring server, use the `get` command with the `netport` option. The default port number is 1100.

Command	GET
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	GET <option> [<i>machine</i>]
Options	<ul style="list-style-type: none">◆ option—See <i>Storage Mirroring Text Client program settings</i> on page 14-15.◆ machine—Name of the machine
Examples	<ul style="list-style-type: none">◆ <code>get netport</code>◆ <code>get VerifyLogName</code>◆ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. To modify the netport setting on a Storage Mirroring server, use the `set` command with the `netport` option. The `netport` option can be any number from 1025-65535 that you want to assign to the port.

```
Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)

=====
Command:
=====
> set netport = 1110
```

3. To determine the current unetport setting, use the `get` command with the `unetport` option. The default port number is 1100.
4. To modify the unetport, use the `set` command with the `unetport` option. The `unetport` option can be any number from 1025-65535 that you want to assign to the port.

Getting Help in the Text Clients

A listing of the DTCL commands and their syntax is available in the Storage Mirroring Text Client and Command Line Client by typing the `help` command.

Command	<code>HELP</code>
Description	Displays the DTCL commands and their syntax
Syntax	<code>HELP</code>
Notes	<ul style="list-style-type: none">◆ Press any key to scroll through the list of commands.◆ Press <code>q</code> to exit the help function.◆ You can also type <code>dtcl help</code> from the directory where the Storage Mirroring program files are installed to display the DTCL commands and their syntax.

Exiting the Text Clients

To exit the Storage Mirroring Text Client and Command Line Client, type the `exit` or `quit` command.

Command	<code>EXIT</code>
Description	Exits the Text Client and the Command Line Interactive client
Syntax	<code>EXIT</code>
Command	<code>QUIT</code>
Description	Quits the Text Client and the Command Line Interactive client
Syntax	<code>QUIT</code>

4 ► Failover Control Center

The Storage Mirroring Failover Control Center is the client that controls all aspects of Storage Mirroring failover. The other clients control different subsets of the failover features.

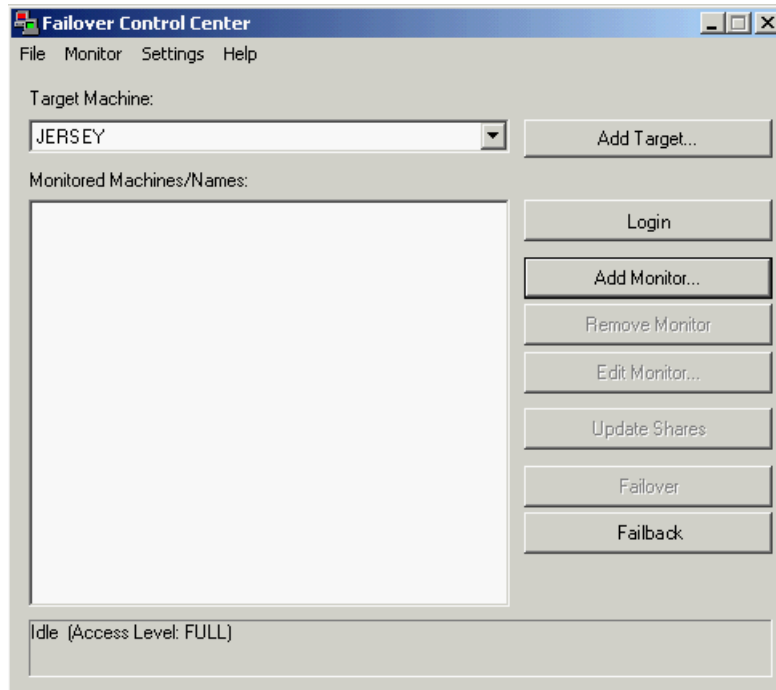
- ◆ **Management Console**—This client can configure the failover settings.
- ◆ **Text Client**—This client can configure the failover settings and can control the failover processes.
- ◆ **Failover Control Center**—This client can configure the failover settings, control the failover processes, and monitor the status of failover activity.

For detailed steps on how failover works including configuring and monitoring failover, see [Failover](#) on page 11-1.

Starting the Failover Control Center

The Failover Control Center can be started from within the Management Console or from the Windows desktop.

- ◆ From the Management Console, select **Tools, Failover Control Center**.
- ◆ From the Windows desktop, select **Start, Programs, Storage Mirroring, Failover Control Center**.



Understanding the Display

The Failover Control Center's main window controls three types of failover activities:

- ◆ Identifying what the target is monitoring
- ◆ Configuring failover operations
- ◆ Monitoring failover activity

Identifying what the target is monitoring

The two main areas of the Control Center screen are the **Target Machine** selection and the **Monitored Machines/Names** pane. When a target machine is selected, any source machines and IP addresses that are currently being monitored are displayed in an expandable and collapsible tree view. Selecting a different target machine will update the tree to reflect the machines and IP addresses that the selected target is monitoring.

Configuring failover operations

The operational buttons located along the right side of the Failover Control Center window access failover configuration dialog boxes and initiate failover processes.


- ◆ **Add Target**—Allows you to specify a target machine that is not available in the **Target Machines** list and log on to it.
- ◆ **Login**—Allows you to log on to a selected machine.
- ◆ **Add Monitor**—Allows you to add a source machine to be monitored by the selected target and log onto it.
- ◆ **Remove Monitor**—Removes a source machine that no longer needs to be monitored.
- ◆ **Edit Monitor**—Allows you to edit the failover settings for a source machine that is currently being monitored by the selected target.
- ◆ **Update Shares**—The target machine manually retrieves information on the existing shares that exist on the source machine.
- ◆ **Failover**—Manually initiates the failover process.
- ◆ **Failback**—Manually initiates the failback process.

For complete details on configuring failover, see [Failover](#) on page 11-1.

Monitoring failover

Now that replication and failover monitoring are configured and started, you will need to know if and when there is a problem. Since it can be essential to quickly know the status of your machines, Storage Mirroring offers various methods for monitoring the status of failover. When the Failover Control Center is running, you will see four visual indicators:

- ◆ The Failover Control Center Time to Fail counter
- ◆ The Failover Control Center status bar located at the bottom of the window
- ◆ The Failover Control Center colored bullets to the left of each IP address and source machine
- ◆ The Windows desktop icon tray containing a failover icon



NOTE: You can minimize the Failover Control Center and, although it will not appear in your Windows taskbar, it will still be active and the failover icon will still appear in the desktop icon tray.

The Failover Control Center does not have to be running for failover to occur.

The following table identifies how the visual indicators change as the status of failover changes.

	Time to Fail Countdown	Status Bar	Colored Bullets	Desktop Icon Tray
Source is Online	The Time to Fail counter is counting down and resetting each time a heartbeat is received from the source machine.	The status bar indicates that the target machine is monitoring the source machine.	The bullets are green. ^a	The Windows desktop icon tray contains a failover icon with red and green computers.
Source Fails and Failover is Initiated	The Time to Fail countdown value is 0.	The status bar displays the source machine and IP address currently being assumed by the target.	The bullets are red.	The Windows desktop icon tray contains a failover icon with red and green computers.
Failover is Complete	The Time to Fail counter is replaced with the "Failed Over" message.	The status bar indicates that monitoring has continued.	The bullets are red.	The Windows desktop icon tray contains a failover icon with a red computer.


a. When the **Time to Fail** value has decreased by 25% of the entire timeout period, the bullet changes from green to yellow, indicating that the target has not received a response from the source. The yellow bullet is a caution signal. If a response from the source is received, the countdown resets and the bullets change back to green. If the countown reaches zero without the target receiving a response from the source, failover begins.

Communication ports

The Storage Mirroring service and the Failover Control Center use two ports for various Storage Mirroring communications.

- ◆ **Service Transmit Port**—The Failover Control Center uses this port to send commands to Storage Mirroring servers.
- ◆ **Heartbeat Listen Port**—The Failover Control Center uses this port to listen for heartbeats transmitted from Storage Mirroring servers.

To view or modify the port settings in the Failover Control Center, select **Settings, Communications**.



NOTE: Since the Management Console and the Failover Control Center share port settings, if the Storage Mirroring service is restarted, changes to the port settings that were made in the Failover Control Center are reflected in the Management Console.

If the Storage Mirroring service is not restarted, the changes will occur in the Failover Control Center only.

5 ► Replication Sets

A replication set defines the data on a source machine that is to be protected by Storage Mirroring. Replication sets are defined by the volumes, directories, files, or wild card combinations that are to be replicated to the target. Creating multiple replication sets allows you to customize sets of data that need to be protected.

Replication Set Rules

Replication sets are created by defining rules. Each replication set rule consists of:

- ◆ **Path**—The path including volume, drive, directory, file, and/or wild card
- ◆ **Include**—If the path and/or file is to be included in the files sent to the target
- ◆ **Exclude**—If the path and/or file is not to be included in the files sent to the target
- ◆ **Recursive**—If the path should automatically be applied to the subdirectories of the named directory

NOTE: Rules can be written for paths that do not currently exist. To enter a nonexistent path, you will have to manually enter the rule.

For example, a replication set rule might be

```
volume\directory\* inc, rec
```

This specifies that all files contained in the volume\directory path are included in the replication set. Because recursion is set, all files and subdirectories under volume\directory are also included. The replication set becomes a list of replication set rules. For example, suppose you want to protect your network server which contains an important customer database. You might create a replication set with the following rule.

```
\exchsrvr\*.db inc, rec
```

This rule includes all database files from the `exchsrvr` directory. Recursion is specified with this rule so all `.db` files in any subdirectory under the `exchsrvr` directory will also be included in the replication set.

NOTE: Replication set rules are limited in length meaning that the entire volume\directory\filename including slashes, spaces, periods, extensions, cannot exceed 259 characters.

Storage Mirroring can mirror, replicate, verify, and restore file names up to 32,760 characters, although each individual component (file or directory name) is limited to 259 characters. File names longer than 32,760 characters (due to the concatenation of the source and target paths) will be skipped and logged to the Storage Mirroring log file and the Windows Event Viewer.

Replication Set Operations

The following replication set operations are available:

- ◆ **Create**—You can create replication sets through the Management Console or the text clients by selecting icons through the graphical interface or by specifying explicit paths through any of the clients.
- ◆ **Modify**—Replication set modifications are a dynamic process. While Storage Mirroring is replicating, changes can be made to a replication set that is in use. Once the changes have been saved, the new settings will become effective immediately. If new folders or files are added to a set that is being replicated, a remirror is necessary to synchronize the source and target.

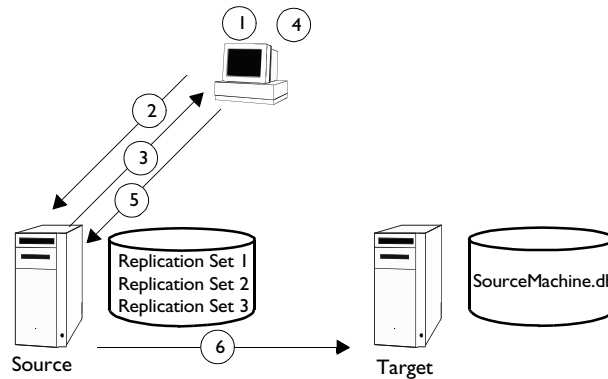
In certain circumstances, the Management Console will display a window stating that a modified replication set needs to be disconnected and reconnected after a save has been completed. This is to allow Storage Mirroring to read the new data that was added to the replication set.

Modifications in the Management Console consist of clearing and selecting alternate volumes, directories, and/or files while modifications in the text clients consist of deleting an existing rule and adding a new one.

- ◆ **Rename/Copy**—You can rename or copy a replication set in the Management Console only. The text clients require you to create a new replication set. When renaming an unsaved replication set, you supply a new name and the original name is replaced. If the replication set has already been saved, a copy is made; you supply a new name and the original is left intact.
- ◆ **Save**—A replication set must be saved before establishing a source/target connection and before exiting any of the Storage Mirroring clients so that it will be available in future client sessions.
- ◆ **Delete**—If a replication set is not currently connected to a target, you can delete that replication set. After a source/target connection has been established, a replication set cannot be deleted until the source and target have been disconnected.
- ◆ **Display/List**—When using the text clients, you can display all of the replication sets associated with a specified source or you can list the rules that are associated with a specified replication set.
- ◆ **Resync/Revert**—You can restore previously saved settings. In this instance, a previously saved replication set where changes were made, but are no longer wanted, can be discarded without losing the replication set altogether. You can revert back to the replication set that was saved before making the changes.

How Replication Sets Work

To better understand how replication sets work, the following diagram identifies the steps involved with creating and saving a replication set as well as what happens to the replication set information after a source and target have been connected. Each numbered step is described following the diagram.



1. The Storage Mirroring client manages all replication set modifications. The administrator creates, modifies, or deletes replication set rules or the replication set itself from the Storage Mirroring client machine.
2. The first time a replication set is saved, the replication set database, `Dbldata.db`, is created on the source machine. This file contains all of the replication sets that are established for that machine.
3. When the Storage Mirroring client manages any replication set activity after the initial creation of `Dbldata.db`, the source machine sends the entire file to the client machine.
4. When the administrator modifies a replication set, the Storage Mirroring client manages that activity and modifies the local copy of `Dbldata.db`.
5. When the replication set is saved, the `Dbldata.db` file is sent to the source machine.
6. When a source/target connection is established, a portion of the replication set database is sent to the target machine. The portion contains only the data that is relevant to the current connection; it does not contain replication set data that is relevant to other connections. The file on the target machine is identical to the source machine's name so that the file can be recovered if the source machine crashes. It is called `source_machine_name.db`.

Multi-Client Access

Replication sets contain error checking to avoid inadvertent overwrites of the `Dbltake.db` file. When `Dbltake.db` is retrieved from the source to the client, a generation number is associated with the retrieval. The generation number is incremented anytime there are successful changes to `Dbltake.db`. When attempting to save a replication set, the generation numbers on the source and client are compared and if they are different, the save is not allowed. The Storage Mirroring Management Console will prompt you if the replication set cannot be saved. This error checking safeguards the replication set data in the event that more than one client machine is accessing the source machine's replication sets.

Replication Capabilities

Storage Mirroring replicates all file and directory data stored on any Windows file system (FAT, FAT32, NTFS4, and NTFS5). Replicated items also include Macintosh[®] files, compressed files, NTFS attributes and ACLs, dynamic volumes, and sparse files. Files can be replicated across mount points, even though mount points are not created on the target.

Storage Mirroring does not replicate items that are not stored on the file system, such as physical volume data and registry based data. By default, system volume information, if included in the replication set, will be mirrored and replicated, although it is recommended that this type of data not be included in the replication set. The %systemroot%, usually c:\winnt, is not mirrored or replicated, by default. Additionally, Storage Mirroring does not replicate NTFS extended attributes, registry hive files, Windows or any system or driver pagefile, system metadata files (\$LogFile, \$Mft, \$BitMap, \$Extend\\$\Usnjrnl, \$Extend\\$\Quota, \$Extend\\$\ObjId, and \$Extend\\$\Reparse), hard links, or the Storage Mirroring disk-based queue logs.

In addition, note the following information:

1. If you select a Windows 200x mount point, by default, the mount point data will be stored in a directory on the target machine. You can create a mount point on the target to store the data or maintain the replicated data in a directory. If you use a directory, it must be able to handle the amount of data contained in the mount point.
2. If you select a Windows 200x dynamic volume and you increase the size of the volume, the target machine must be able to compensate for an increase in the size of the dynamic volume.
3. If you select data stored on a recursive mount point for replication, a mirror will never finish. Storage Mirroring does not check for data stored on recursive mount points.
4. Reparse points are not replicated by Storage Mirroring.
5. If you are using alternate data streams, they are not included in the replication set size calculation. Therefore, you may see the mirror process at 100% complete while mirroring continues.
6. If you select Windows 200x encrypted files for replication, keep in mind the following:
 - a. Only the data, not the attributes or security/ownership, is replicated. However, the encryption key is included. This means that only the person who created the encrypted file on the source will have access to it on the target.
 - b. Only data changes cause replication to occur; changing security/ownership or attributes does not.
 - c. Replication will not occur until the Windows Cache Manager has released the file. This may take awhile, but replication will occur when Storage Mirroring can access the file.
 - d. When remirroring, the entire file is transmitted every time, regardless of the remirror settings.
 - e. Verification cannot check encrypted files because of the encryption. If remirror is selected, the entire encrypted file will be remirrored to the target. Independent of the remirror option, all encrypted files will be identified in the verification log.
7. If you have a FAT volume on the source, mounted on a directory which resides on an NTFS volume, also on the source, these files will not be mirrored, regardless of the target file system. Replication will work correctly. To work around this issue, make sure both volumes are NTFS.
8. If you select a compressed file or folder from an NTFS partition and replicate it to FAT target, the attributes are lost but the data is maintained.
9. If you are mirroring/replicating from an NTFS source to a FAT target, you may see additional error messages in your Storage Mirroring log file because the target file system cannot handle the NTFS attributes or file permissions. Here are two examples:
 - a. Alternate Data Streams—If your replication set contains files with alternate data streams, you will see messages indicating that there are unfinished operations because the FAT file system cannot store the alternate, or the default, data stream information.
 - b. Encrypted files—If your replication set contains encrypted files, the files will not be replicated. They will exist on the target as zero byte files (0 KB) and the Storage Mirroring log will report that the files are not backed up.

NOTE: If any directory or file contained in your replication set specifically denies permission to the system account or the account running the Storage Mirroring service, the attributes of the file on the target will not be updated because of the lack of access. This also includes denying permission to the Everyone group because this group contains the system account.

Managing Replication Sets Through the Management Console

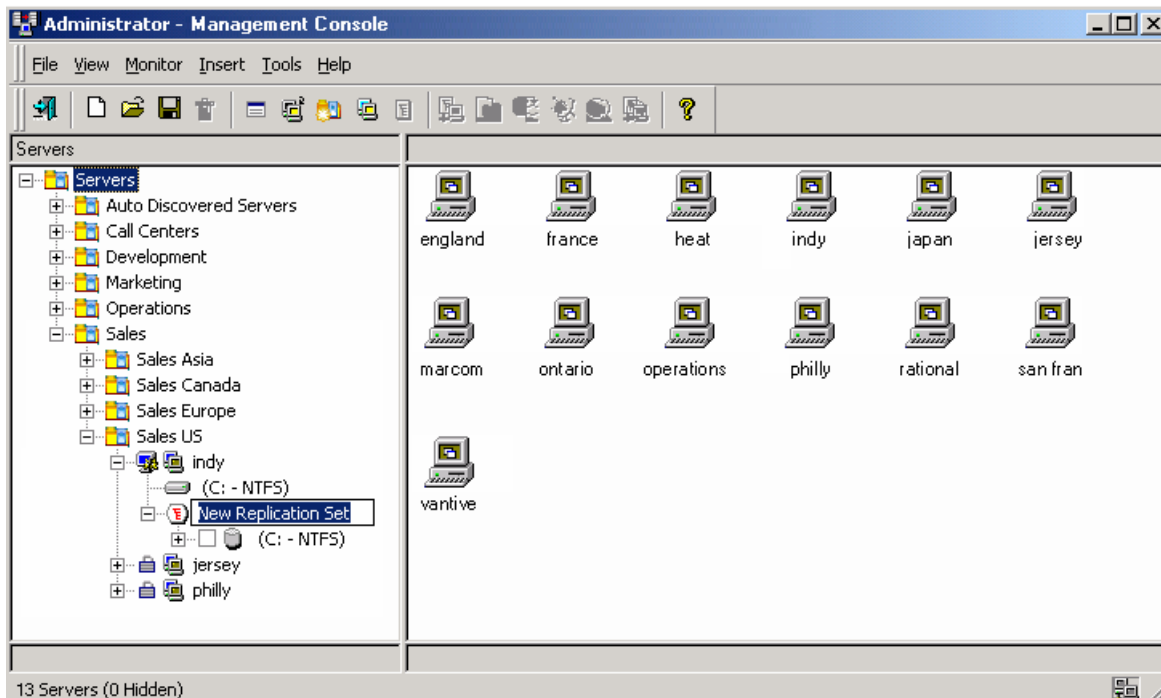
The following sections describe how to manage replication sets through the Storage Mirroring Management Console. Each section assumes that you have already logged into the source machine.

Creating a replication set

1. You can use either of the following methods to create a new replication set when a source machine is highlighted:

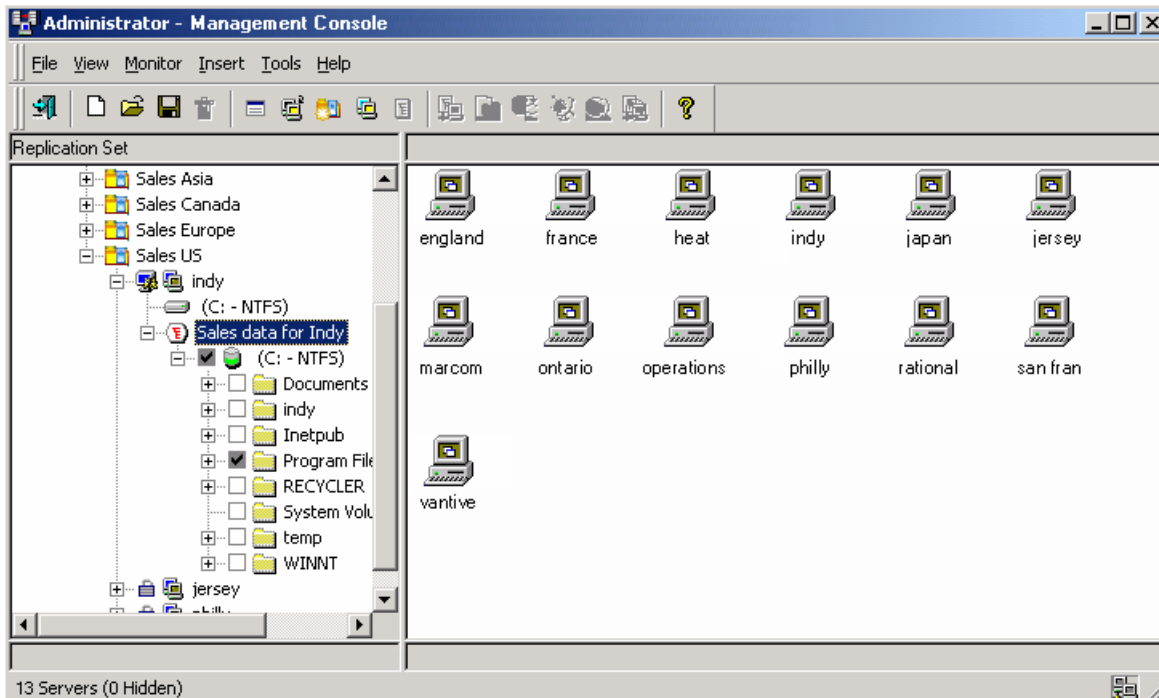
- ◆ Select **Insert, Replication Set** from the menu bar.
- ◆ Right-click on the source machine name and select **New, Replication Set**.

A replication set icon appears in the left pane. By default, it is named `New Replication Set`.



2. Rename the newly inserted replication set with a unique name by typing over the default name and pressing **Enter**. This process is similar to naming a new folder in Windows Explorer.

3. To see the directory tree for the source machine selected, expand and collapse the directory structure in the left pane of the Management Console by clicking on a box containing a plus sign to expand the tree and a box containing a minus sign to collapse the tree.



NOTE: The default number of files that are listed in the right pane of the Management Console is 2500, but this is user configurable. A larger number of file listings allows you to see more files in the Management Console, but results in a slower display rate. A smaller number of file listings displays faster, but may not show all files contained in the directory. To change the number of files displayed, select **File, Options** and adjust the **File Listings** sidebar to the desired number.

4. Define the data to be protected by selecting volumes, drives, directories, and/or specific files by marking the check box next to the desired selection.

NOTE: Be sure and verify what files can be included by reviewing [Replication Capabilities](#) on page 5-6.


5. After selecting the data for this replication set, right-click the new replication set icon and select **Save**. A saved replication set icon will change from red to black.

NOTE: To connect this replication set to a target machine, continue with [Establishing a new connection](#) on page 6-8.

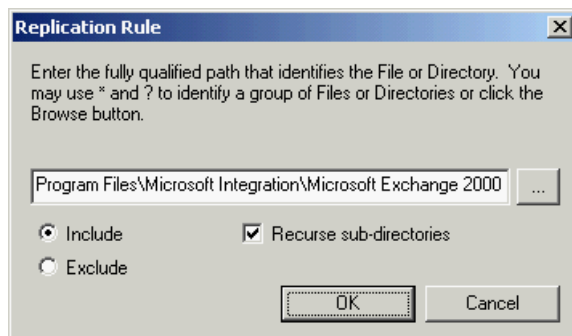
Creating a replication set using wild cards and direct rule entry

In some instances, data may need to be replicated that cannot be easily selected from the Management Console, or the directories or files may not exist. For example, the Storage Mirroring wild card capabilities and flexible file selection process allows you to manually include all .db files from a specific directory that is otherwise excluded. This flexibility is available through the Replication Set Properties dialog box.

1. Right-click on the replication set icon and select **Properties**. The Replication Set Properties dialog box appears and lists any rules that have been defined from the main Management Console window. Each definition will also display the attributes associated with that replication set rule:
 - ◆ **Inc**—Include indicates that the directory or wild card specifications are a part of the replication set. These files will be sent to the target machine.
 - ◆ **Exc**—Exclude indicates that the directory or wild card specifications are not a part of the replication set. These files will not be sent to the target machine.
 - ◆ **Rec**—Recursion indicates that the directory or wild card specifications are selected recursively. All files and subdirectories under the defined path will be sent to the target machine.



2. From the Replication Set Properties dialog box, click **Add**.
3. Specify a path and, if necessary, the wild card or specific file name. Select the **Include**, **Exclude**, and/or **Recurse sub-directories** attributes to be applied to this definition and click **OK** to complete the definition of the rule.




4. Repeat steps 2 and 3 to add additional replication set rules.
5. After all additional replication set rules have been defined, exit the Replication Set Properties dialog box by clicking **OK** or click **Cancel** to exit the dialog box without saving the changes.
6. After defining the rules for this replication set, right-click the replication set icon and select **Save**. A saved replication set icon will change from red to black.

Calculating replication set size

While Storage Mirroring is mirroring, the right pane of the Management Console displays statistics to keep you informed of its progress. If the size of the replication set is determined before the mirror is started, Storage Mirroring can display the percentage of the replication set that has been mirrored in the **Mirror Status** column. If the size was not calculated prior to starting the mirror, the column displays **Mirroring**. You can calculate the replication set size through the Replication Set Properties dialog box by using the instructions below or through the Connection Manager (see [Initiating a mirror upon connection](#) on page 7-6).

1. Right-click on the replication set icon and select **Properties**. The Replication Set Properties dialog box appears.
2. Select the **Size** tab.
3. If the replication set size has never been determined, click **Calculate**. If the replication set has previously been determined, the button will be labeled **Recalculate**.



NOTE: Depending on user activity, the size shown may not accurately reflect the current size of the replication set. If changes are occurring to files in the replication set while the calculation is being made, the actual size may differ slightly. The amount of data is determined at the exact time the calculation is made.

4. Click **OK** to return to the Management Console.

Modifying a replication set

Storage Mirroring allows you to make modifications to a replication set when you want to change the data you wish to protect. This allows you to add, remove, or modify any replication set rules without having to create a new replication set.

1. In the left pane, highlight the replication set you want to modify and expand the volume and directory levels as needed.
2. Modify any items defined through the main Management Console window by marking or clearing the volume, drive, directory, or file check boxes.
3. Modify any items defined through the Replication Set Properties dialog box by right-clicking the replication set name and selecting **Properties**.
 - a. Edit an existing rule by highlighting it, selecting **Edit**, changing the **Include**, **Exclude**, or **Recurse sub-directories** attributes; defining a new path; or specifying a new file or wild card. Click **OK** when the modifications are complete.
 - b. Remove a rule by highlighting it and selecting **Remove**.
 - c. Add a rule as described in [Creating a replication set using wild cards and direct rule entry](#) on page 5-9.
4. After updating the rules for this replication set, right-click the new replication set icon and select **Save**. A saved replication set icon will change from red to black.

NOTE: If you save changes to a connected replication set, it is recommended that you perform a mirror to guarantee data integrity between the source and target machines. A dialog box will appear allowing you to choose to initiate a mirror immediately. If you are using Storage Mirroring to provide disaster recovery for Microsoft Cluster Server, you must modify the replication set on both nodes. The replication set must be the same on both nodes for Storage Mirroring to function properly. For detailed information, see [Providing Disaster Recovery Using Storage Mirroring](#) on page C-4.

Renaming and copying replication sets

To rename or copy a replication set, left-click once on a highlighted replication set name to edit the field. Specify a unique name and press **Enter**. The process is similar to renaming a folder in Windows Explorer. If the original replication set has not been saved, the new name replaces the original name. If the original replication set is saved, the new name creates a copy of the original replication set.

Deleting a replication set

To delete a replication set, right-click the replication set icon and select **Delete** or press the **Delete** key on the keyboard.

NOTE: If a replication set is not currently connected to a target, you can delete that replication set. After a source/target connection has been established, a replication set cannot be deleted until the source and target have been disconnected. To disconnect a source and target, see [Disconnecting a connection](#) on page 6-9.

If you are using Storage Mirroring to provide disaster recovery for Microsoft Cluster Server, you must modify the replication set on both nodes. The replication set must be the same on both nodes for Storage Mirroring to function properly. For detailed information, see [Providing Disaster Recovery Using Storage Mirroring](#) on page C-4.

Managing Replication Sets Through the Text Client

The following sections describe how to manage replication sets through the Storage Mirroring text clients. For a complete list of the DTCL commands, see [Storage Mirroring Commands](#) on page A-1. Each of the following commands assumes that you have already logged on and identified a source machine.

Creating a replication set

1. Use the `repset create` command to create a new replication set.

Command	<code>REPSET CREATE</code>
Description	Creates a replication set
Syntax	<code>REPSET CREATE <name></code>
Options	<code>name</code> —Name of the replication set
Examples	<ul style="list-style-type: none">◆ <code>repset create "Exchange Repset"</code>◆ <code>repset create marketing</code>
Notes	<ul style="list-style-type: none">◆ If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.◆ If the name of the replication set is the same as a word used in a DTCL command (for example, <code>repset create repset</code>), you will receive an error. If you want to use a word like this, you must enclose it in quotation marks (<code>repset create "repset"</code>).

2. Specify it as the active replication set by using the `repset use` command.

Command	<code>REPSET USE</code>
Description	Specifies a replication set as the active replication set
Syntax	<code>REPSET USE <repset></code>
Options	<code>repset</code> —Name of the replication set
Examples	<ul style="list-style-type: none">◆ <code>repset use "Exchange Repset"</code>◆ <code>repset use marketing_database</code>
Notes	If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.

```
Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)

=====
Command:
=====
repset create exchange
> repset use exchange
```


-
3. Define a replication set rule by using the `repset rule add` command.

Command	<code>REPSET RULE ADD</code>
Description	Adds a rule to a replication set. A rule is the specification of a path including volume, directories, wild cards, and/or file names.
Syntax	<code>REPSET RULE ADD <path> [<u>I</u>NCLUDE <u>E</u>XCLUDE][, <u>R</u>ECURSIVE <u>N</u>ONRECURSIVE] [TO <repset>]</code>
Options	<ul style="list-style-type: none">◆ <code>path</code>—Volume, directory, wild card, and/or file name◆ <code>INCLUDE</code>—Include the specified path in the replication set◆ <code>EXCLUDE</code>—Exclude the specified path in the replication set◆ <code>RECURSIVE</code>—All subdirectories and files are recursively included or excluded◆ <code>NONRECURSIVE</code>—All subdirectories and files are non-recursively included or excluded◆ <code>repset</code>—Name of the replication set
Examples	<code>repset rule add c:\exchange to "Exchange Repset"</code>
Notes	<ul style="list-style-type: none">◆ The default settings for this command are include and recursive.◆ If you do not specify a replication set name, the current replication set will be used.◆ If the path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.◆ Verify what files can be included by reviewing Multi-Client Access on page 5-5.

4. Repeat the `repset rule add` command to completely define your replication set.
5. If you need to remove a rule, use the `repset rule remove` command.

Command	<code>REPSET RULE REMOVE</code>
Description	Removes a rule from a replication set
Syntax	<code>REPSET RULE <u>R</u>EMOVE <path> [<u>F</u>ROM <repset>]</code>
Options	<ul style="list-style-type: none">◆ <code>path</code>—Volume, directory, wild card, and/or file name◆ <code>repset</code>—Name of the replication set
Examples	<code>repset rule remove c:\exchange from "Exchange Repset"</code>
Notes	<ul style="list-style-type: none">◆ If you do not specify a replication set name, the current replication set will be used.◆ If the path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.

6. After you have added all of the rules, save the replication set by using the `repset save` command.

Command	<code>REPSET SAVE</code>
Description	Saves all replication set rules for the currently selected source
Syntax	<code>REPSET SAVE</code>

NOTE: To connect this replication set to a target machine, continue with [Establishing a new connection](#) on page 6-16.

Calculating replication set size

While Storage Mirroring is mirroring, the Text Client displays statistics to keep you informed of its progress. If the size of the replication set is determined before the mirror is started, Storage Mirroring can display the percentage of the replication set that has been mirrored. If the size was not calculated prior to starting the mirror, it displays **Mirroring**. You can calculate the replication set size manually or on connection.

1. Use the `repset calculate` DTCL command to manually calculate the size of the replication set.

Command	<code>REPSET CALCULATE</code>
----------------	-------------------------------

Description	Calculates the size of a replication set
--------------------	--

Syntax	<code>REPSET <u>CALCULATE</u> [<i>repset</i>]</code>
---------------	--

Options	<i>repset</i> —Name of the replication set
----------------	--

Examples	<ul style="list-style-type: none">◆ <code>repset calculate marketing</code>◆ <code>repset calc "Exchange repset"</code>
-----------------	--

Notes	<ul style="list-style-type: none">◆ If a replication set name is not specified, the active replication set will be used.◆ If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.◆ The results of the calculation are logged to the Storage Mirroring log file.
--------------	---

2. To calculate the size of the replication set automatically upon connection, verify the setting of the `CalculateOnConnect` option.
 - a. To determine if the replication set size will be calculated when the connection is established, use the `get` command with the `CalculateOnConnect` option.

Command	<code>GET</code>
----------------	------------------

Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
--------------------	--

Syntax	<code>GET <<i>option</i>> [<i>machine</i>]</code>
---------------	---

Options	<ul style="list-style-type: none">◆ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.◆ <i>machine</i>—Name of the machine
----------------	--

Examples	<ul style="list-style-type: none">◆ <code>get netport</code>◆ <code>get VerifyLogName</code>◆ <code>get DefaultAddress</code>
-----------------	---

Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.
--------------	--

-
- b. To modify whether or not the replication set size will be calculated when the connection is established, use the `set` command with the `CalculateOnConnect` option. The `CalculateOnConnect` option is set to 1 for enabled or 0 for disabled. By default, the replication set size will be calculated.

Command	<code>SET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	<code>SET <option>=<value> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>set netport=1100</code>♦ <code>set VerifyLogName="HA server.log"</code>♦ <code>set DefaultAddress="129.12.34.3"</code>
Notes	<ul style="list-style-type: none">♦ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed.♦ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

Modifying a replication set

Make modifications to a replication set when you want to change the data you wish to protect.

1. View the replication sets associated with the active source machine by using the `repset list` command.

Command	<code>REPSET LIST</code>
Description	Lists all replication set names for the currently selected source
Syntax	<code>REPSET LIST</code>

```
Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)

- List of rep sets -      * DATABASE HAS BEEN MODIFIED *
exchange                  disabled
  (press any key)

> repset list
```

2. Identify a replication set as active by using the `repset use` command.

Command	<code>REPSET USE</code>
Description	Specifies a replication set as the active replication set
Syntax	<code>REPSET USE <repset></code>
Options	<code>repset</code> —Name of the replication set
Examples	<ul style="list-style-type: none">◆ <code>repset use "Exchange Repset"</code>◆ <code>repset use marketing_database</code>
Notes	If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.

3. View the replication set's rules by using the `repset display` command.

Command	<code>REPSET DISPLAY</code>
Description	Displays the rule of a replication set
Syntax	<code>REPSET <u>DISPLAY</u> [<code>repset</code>]</code>
Options	<code>repset</code> —Name of the replication set
Examples	<ul style="list-style-type: none">◆ <code>repset display</code>◆ <code>repset display Exchange</code>
Notes	<ul style="list-style-type: none">◆ If you do not specify a replication set name, the current replication set will be used.◆ If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.

-
4. Remove the existing rule that you need to change using the `repset rule remove` command.

Command	<code>REPSET RULE REMOVE</code>
Description	Removes a rule from a replication set
Syntax	<code>REPSET RULE <u>REMOVE</u> <path> [FROM <repset>]</code>
Options	<ul style="list-style-type: none">◆ path—Volume, directory, wild card, and/or file name◆ repset—Name of the replication set
Examples	<code>repset rule remove c:\exchange from "Exchange Repset"</code>
Notes	<ul style="list-style-type: none">◆ If you do not specify a replication set name, the current replication set will be used.◆ If the path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.

5. Add a new rule by using the `repset rule add` command.

Command	<code>REPSET RULE ADD</code>
Description	Adds a rule to a replication set. A rule is the specification of a path including volume, directories, wild cards, and/or file names.
Syntax	<code>REPSET RULE ADD <path> [<u>INCLUDE</u> <u>EXCLUDE</u>][, <u>RECURSIVE</u> <u>NONRECURSIVE</u>] [TO <repset>]</code>
Options	<ul style="list-style-type: none">◆ path—Volume, directory, wild card, and/or file name◆ INCLUDE—Include the specified path in the replication set◆ EXCLUDE—Exclude the specified path in the replication set◆ RECURSIVE—All subdirectories and files are recursively included or excluded◆ NONRECURSIVE—All subdirectories and files are non-recursively included or excluded◆ repset—Name of the replication set
Examples	<code>repset rule add c:\exchange to "Exchange Repset"</code>
Notes	<ul style="list-style-type: none">◆ The default settings for this command are include and recursive.◆ If you do not specify a replication set name, the current replication set will be used.◆ If the path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.◆ Verify what files can be included by reviewing Multi-Client Access on page 5-5.

-
6. If you decide that you do not like the changes you have just made, you can use the `repset resync` command to restore the previously saved settings, undoing your modifications.

NOTE: You will lose any changes that you have made if you use the `repset resync` command.

Command `REPSET RESYNC`

Description Retrieves the last saved replication set settings from the source.

Syntax `REPSET RESYNC`

7. Save the replication set by using the `repset save` command.

Command `REPSET SAVE`

Description Saves all replication set rules for the currently selected source

Syntax `REPSET SAVE`

Deleting a replication set

1. View the replication sets associated with the active source machine by using the `repset list` command.

Command	<code>REPSET LIST</code>
Description	Lists all replication set names for the currently selected source
Syntax	<code>REPSET LIST</code>

```
Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)

- List of rep sets -      * DATABASE HAS BEEN MODIFIED *
exchange                  disabled
  (press any key)

> repset list
```

2. Delete the replication set by using the `repset delete` command.

Command	<code>REPSET DELETE</code>
Description	Deletes the specified replication set
Syntax	<code>REPSET <u>DELETE</u> <repset></code>
Options	<code>repset</code> —Name of the replication set
Examples	<code>repset delete "Exchange Repset"</code>
Notes	If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.

3. After deleting the replication set, use the `repset save` command so that the deletion will be registered by other Storage Mirroring clients.

Command	<code>REPSET SAVE</code>
Description	Saves all replication set rules for the currently selected source
Syntax	<code>REPSET SAVE</code>

6 ► Connections

The Storage Mirroring connection is the link between the source replication set and a target machine. It is a logical connection and refers to the relationship between the replication set on the source and the copy of the data on the target, as well as the status of the connection. The different connection states are described below.

Connection Status	Description
Started	The network connection exists and is available for data transmission. Replication and mirror data are transmitted to the target as soon as possible. This is the standard state that you will see most often.
Stopped	The source and target machines have been linked through the Storage Mirroring client software, but the network connection does not exist. Replication and mirror data are not transmitted to the target, but are held in queue on the source.
Paused	The network connection exists and is available for data transmission, but the replication and mirror data is being held in a queue and is not being transmitted to the target.
Scheduled	The source and target machines have been linked through the Storage Mirroring client software, but the network connection is not established until event driven or scheduling criteria have been met.
Error	A transmission error has occurred. Possible errors include a broken physical line or a failed target service.

NOTE: These connection states are under the **Transmit Mode** column in the Management Console and the **Transmit** column in the Text Client.

Connection Options

When establishing a connection, you must specify the machines that are involved with the connection, what data is being protected and where that data is going, as well as configuring how this connection will perform.

- ◆ **Source Server**—You must specify the Storage Mirroring source machine that contains the replication set that is going to be transmitted to the Storage Mirroring target.
- ◆ **Replication Set**—At least one replication set must exist on the source before establishing a connection. You must specify the replication set that will be connected to the target machine.

NOTE: The same replication set can be connected to multiple target machines.

- ◆ **Target Server**—You must specify which Storage Mirroring target will maintain the copy of the source data. You can specify a machine name, IP address, or virtual IP address.
- ◆ **Route**—This is an optional setting allowing you to specify the IP address on the target the data will be transmitted through. You should only change this setting if you want to select a different route for Storage Mirroring traffic. On a machine with more than one NIC, this increases the flexibility of configuring Storage Mirroring activity. For example, you can separate regular network traffic and Storage Mirroring traffic on a machine.
- ◆ **Mappings**—You must specify the location on the target where the replicated data will reside. Storage Mirroring offers two predefined locations as well as a custom option that allows you to create your own path.

- ◆ **One-to-One**—In cases where only one source is replicated to each target, it is often easier to maintain the identical volume/directory structure on the source and target. Additionally, if the target will stand-in for the source during failover where users will be accessing data directly on the target (versus restoring data to the source for user access or accessing data through shares or exports), this configuration is generally used.

The one-to-one option in the Management Console and the `map exact` switch used with the `connect` command in the Text Client map each source volume to the same structure on the target. For example, `c:\data` and `d:\files` on the source would be replicated to `c:\data` and `d:\files`, respectively, on the target.

- ◆ **All-to-One**—In cases where more than one source is replicated to a single target server, it may be necessary to store the data from each source in a separate directory on the same volume. If separate target locations are not defined for each source, there is a potential that replicated data could be overwritten and corrupted. Replicating each source to its own directory structure also simplifies the restore process.

The all-to-one option in the Management Console and the `map base` switch used with the `connect` command in the Text Client, map each source volume to a single target volume.

The predefined path is `\source_name\repset_name\ volume_name`. For example, the default path for the replication set rule `c:\data` is `\source_name\ replication_set_name\c`.

- ◆ **Custom Location**—If the predefined options do not store the data in a location that is appropriate for your network operations, you can specify your own custom location where the replicated files will be sent.

NOTE: If you are using Windows 200x, your mapping selection can be critical.

- ◆ If you are mirroring and replicating dynamic volumes or mount points, your location on the target must be able to accommodate the amount of data that you are replicating.
 - ◆ If you are mirroring and replicating sparse files and your location on the target is a non-NTFS 5 volume, the amount of disk space available must be equal to or greater than the entire size of the sparse file. If you are mirroring and replicating to an NTFS 5 volume, the amount of disk space available must be equal to or greater than the on-disc size of the sparse file.
 - ◆ If you are mirroring and replicating multiple mount points, your directory mapping must not create a cycle or loop. For example, if you have the `c:` volume mounted at `d:\c` and the `d:` volume mounted at `c:\d`, this is a circular configuration. If you create and connect a replication set for either `c:\d` or `d:\c`, there will be a circular configuration and mirroring will never complete.
-

-
- ◆ **Start Mirror on Connection**—Mirroring can be initiated immediately when the connection is established. If mirroring is not configured to start automatically, you must start it manually after the connection is established or set a schedule to start it later.

WARNING: Data integrity cannot be guaranteed without a mirror being performed. This option is recommended for the initial connection.

- ◆ **Start Replication on Connection**—Replication can be initiated immediately when the connection is established. If replication is not configured to start automatically, you must start it manually after the connection is established or set a schedule to start it later.

NOTE: If you disable **Start Replication on Connection**, you will need to perform a mirror prior to beginning replication to guarantee data integrity.

Throughput Diagnostic Utility

Built into Storage Mirroring is a null, or non-existent, target that allows you to imitate a normal connection without passing any data across the network. The purpose of this connection is to generate statistics that can be used in approximating the time and amount of bandwidth that this connection will use when actively established. The name of the target is TDU or Throughput Diagnostics Utility. The statistics are logged to a file. By default, the file is called `statistic.sts`. For more information on the statistics logging file, see [DTStat](#) on page 13-30.

Connection IDs

A unique connection ID is associated with each Storage Mirroring connection. The connection ID provides a reference point for each connection and is used as an option in many of the DTCL commands.

Each connection is assigned an ID which is a sequential number starting at one. Each time a connection is established, the ID counter is incremented. It is reset back to one each time the Storage Mirroring service is restarted. For example, if the Storage Mirroring service was started and the same replication set was connected to five target machines, each connection would have a unique connection ID from 1 to 5.

The connection IDs are displayed in the Management Console and the Text Client, but are not displayed in the Command Line Client. If you are using the Command Line Client, you will need to keep track of the connection ID manually or establish a variable so that the connection ID can be used with the DTCL commands. For detailed steps on creating variables, see [Variables](#) on page A-45.

Connection Resources


If you are using Storage Mirroring on a Microsoft Cluster Server, Storage Mirroring has installed a resource to allow you to use a node as a Storage Mirroring source or target. The resource is called the **Double-Take Source Connection**. For complete details on using Storage Mirroring with Microsoft Cluster Server, see [Disaster Recovery for MSCS](#) on page C-1.

Compression

To help reduce the amount of bandwidth needed to transmit Storage Mirroring data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it uncompresses it and then writes it to disk. On a default Storage Mirroring installation, compression is disabled.

How Connections Are Established

To better understand how Storage Mirroring connections work, the following diagram identifies the sequential steps completed when establishing a Storage Mirroring connection. Each numbered step is described below the diagram.




1. When the client initiates a connection, the Storage Mirroring source machine validates the data contained in the replication set and the target information entered in the Connection Manager. The source machine verifies that if the same source and target are connected through a different replication set, the specified target path is not already in use. If the target path is in use by another Storage Mirroring connection, the potential for data overwrites exists. Therefore, if one-to-one configuration is being used, Storage Mirroring will not allow a connection to the same path.
2. After the connection information is validated, the replication set database is sent to the target machine as a backup. The file on the target machine is identical to the source's machine name so that the file can be recovered if the source machine crashes.
3. The source establishes transmission queues. Mirror and replication data is put onto these queues. For example, when using a slow network or target machine, the data is updated on the source faster than it can be transmitted to the target. The data is stored on a queue so that it is not lost.
4. A logical Storage Mirroring connection is established. This is the connection that is visible through the Storage Mirroring client applications.
5. If mirroring or replication begins on connection, a network connection is established and the transmission of data begins immediately. If no activity is set to occur on connection, the network link is established when mirroring and replication are triggered.

Queuing

During the Storage Mirroring installation, you identified the amount of disk space that can be used for Storage Mirroring queuing. Queuing to disk allows Storage Mirroring to accommodate high volume processing that might otherwise fill up system memory. For example, on the source, this may occur if the data is changing faster than it can be transmitted to the target or on the target, a locked file might cause processing to backup.

The following diagram will help you understand how queuing works. Each numbered step is described after the diagram.



1. If data cannot immediately be transmitted to the target, it is stored, or queued, in system memory. You can configure how much system memory you want to use for queuing, but, 128 MB of memory is used by default.
2. When the allocated amount of system memory is full, new changed data bypasses the full system memory and is queued directly to disk. Data queued to disk is written to a transaction log. Each transaction log can store 5 MB worth of data. Once the limit has been reached, a new transaction log is created. The logs can be distinguished by the file name which includes the target IP address, the Storage Mirroring port, the connection ID, and an incrementing sequence number.

NOTE: You may notice transaction log files that are not the defined size limit. This is because data operations are not split. For example, if a transaction log has 10 KB left until the limit and the next operation to be applied to that file is greater than 10 KB, a new transaction log file will be created to store that next operation. Also, if one operation is larger than the defined size limit, the entire operation will be written to one transaction log.

3. When system memory is full, the most recent changed data is added to the disk queue, as described in step 2. This means that system memory contains the oldest data. Therefore, when data is transmitted to the target, Storage Mirroring pulls the data from system memory and sends it. This ensures that the data is transmitted to the target in the same order it was changed on the source. Storage Mirroring automatically reads operations from the oldest transaction log file into system memory. As a transaction log is depleted, it is deleted. When all of the transaction log files are deleted, data is again written directly to system memory (step 1).
4. To ensure the integrity of the data on the target, the information must be applied in the same order as it was on the source. If there are any delays in processing, for example because of a locked file, a similar queuing process occurs on the target. Data that cannot immediately be applied is queued to system memory. By default, 128 MB of system memory is also used on the target.
5. When the allocated amount of system memory on the target is full, new incoming data bypasses the full system memory and is queued directly to disk. Data queued to disk is written to a transaction log. On the target, the transaction logs are identified with the source IP address, the Storage Mirroring port, the connection ID, and an incrementing sequence number.
6. Like the source, system memory on the target contains the oldest data so when data is applied to the target, Storage Mirroring pulls the data from system memory. Storage Mirroring automatically moves operations from the oldest transaction log file to system memory. As a transaction log is depleted, it is deleted. When all of the transaction log files are deleted, data is again written directly to system memory (step 4).

Auto-Disconnect and Auto-Reconnect

While disk queues are user configurable and can be extensive, they are limited by the amount of physical disk space available. If the amount of disk space specified for disk queuing is met, additional data could not be added to the queue and data would be lost. To avoid any data loss, the auto-disconnect and auto-reconnect processes occur.

Exhausted queues on the source—If disk queuing is exhausted on the source, Storage Mirroring will automatically start disconnecting connections. This is called auto-disconnect. The transaction logs and system memory are flushed allowing Storage Mirroring to begin processing anew. The auto-reconnect process ensures that any connections that were auto-disconnected are automatically reconnected. Then, if configured, Storage Mirroring will automatically remirror the data. This process is called auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.

Exhausted queues on the target—If disk queuing is exhausted on the target, the target instructs the source to pause. The source will automatically stop transmitting data to the target and will queue the data changes. When the target recovers, it will automatically tell the source to resume sending data. If the target does not recover by the time the source queues are exhausted, the source will auto-disconnect as described above. The transaction logs and system memory from the source will be flushed then Storage Mirroring will auto-reconnect. If configured, Storage Mirroring will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.

Queuing errors—If there are errors during disk queuing on either the source or target, for example, Storage Mirroring cannot read from or write to the transaction log file, the data integrity cannot be guaranteed. To prevent any loss of data, the source will auto-disconnect and auto-reconnect. If configured, Storage Mirroring will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.

Target server interruption—If a target machine experiences an interruption (such as a cable or NIC failure), the source/target network connection is physically broken but both the source and target maintain the connection information. The Storage Mirroring source, not being able to communicate with the Storage Mirroring target, stops transmitting data to the target and queues the data changes, similar to the exhausted target queues described above. When the interruption is resolved and the physical source/target connection is reestablished, the source begins sending the queued data to the target. If the source/target connection is not reestablished by the time the source queues are exhausted, the source will auto-disconnect as described above.

Target server failure—If a target server is restarted, there could have been data in the target queue when the server was stopped. To prevent any loss of data, the Storage Mirroring service will attempt to persist to disk important target connection information (such as the source and target IP addresses for the connection, various target queue information, the last acknowledged operation, data in memory moved to disk, and so on) before the target server is stopped. If Storage Mirroring is able to successfully persist this information, when the Storage Mirroring service on the target is restarted, Storage Mirroring will pick up where it left off, without requiring an auto-disconnect, auto-reconnect, or auto-remirror. If Storage Mirroring cannot successfully persist this information prior to the restart (for example, a server crash or power failure where the target cannot shutdown gracefully), the source will auto-reconnect when the target is available, and if configured, Storage Mirroring will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.

Managing Connections Through the Management Console

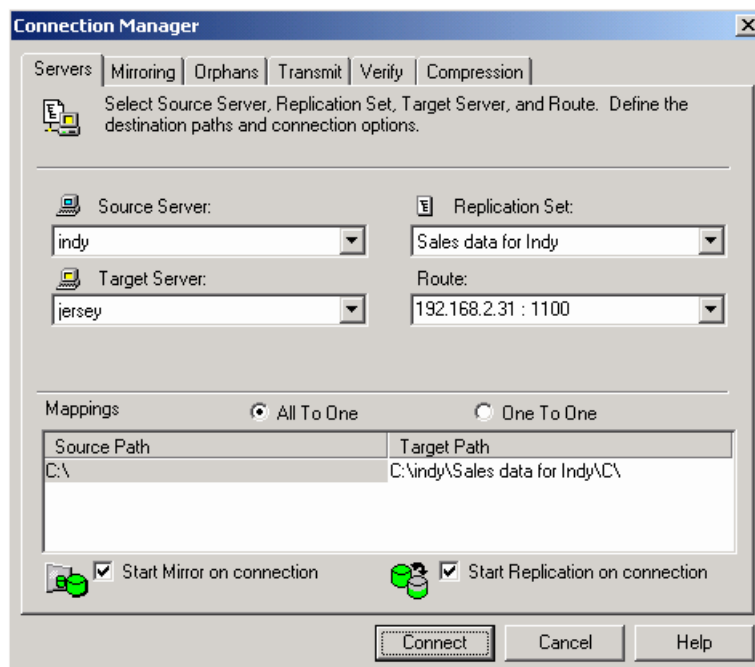
The Management Console can be used to establish, disconnect, and monitor a connection between the source replication set and a target machine. In addition, the Management Console can be used to configure queuing and auto-reconnect.

Establishing a new connection

You can establish a new connection manually through the Connection Manager using the instructions below or using the automated Storage Mirroring Connection Wizard, which prompts you for each option necessary to establish a connection.

1. There are four methods available for opening the Connection Manager and establishing a connection manually:
 - ◆ Highlight a replication set and select **Tools, Connection Manager**.
 - ◆ Right-click on a replication set and select **Connection Manager**.
 - ◆ Drag and drop a replication set to a target machine on the left pane of the Management Console.
 - ◆ Highlight a replication set and then drag and drop the replication set to a target machine on the right pane of the Management Console.

The Connection Manager opens to the **Servers** tab.



NOTE: The Orphans tab is dependent on the target selected, therefore, it will not be displayed until a target is selected.

2. Some entries on the **Servers** tab will be completed depending on which method you used to access it. For example, if you entered the Connection Manager by right-clicking on a replication set, the name of the replication set will be displayed in the replication set field.
3. Verify that the correct source machine and replication set were selected from the Management Console. If they are not, use the **Source Server** and **Replication Set** fields to identify the source machine and replication set for which you want to establish the connection.
4. Depending on the method you used to open the Connection Manager, the **Target Server** and **Route** fields may or may not be completed. If they are not, select a target machine and the primary IP address of that machine will automatically appear. If you have multiple IP addresses on your target, verify that the **Route** field is set to the correct network path.

NOTE: To initiate a connection to the Throughput Diagnostics Utility, select the **Diagnostics** target. The **Route** field will automatically populate with **Throughput Diagnostics Utility (TDU)**. The statistics generated by the TDU can be viewed using DTStat. By default, the file name is `statistic.sts`. For detailed information on statistics logging, see [DTStat](#) on page 13-30.

If you are establishing a connection within a NAT environment, you will need to specify the port of the router after the IP address (seperated by a colon). See [NAT Environments](#) on page D-1 for more details.

-
5. Select the target path, which corresponds to the desired location on the target where the replicated data will reside, by selecting **One-to-One** or **All-to-One**. If you want to specify a custom location, select either of the radio buttons and then click on the directory entry under the Target Path column. You will be in edit mode and can specify the desired location on the target for the replicated data.
 6. If you want mirroring and/or replication to start immediately when the connection is established, mark either or both of the two check boxes at the bottom of the **Servers** tab.

NOTE: Other tabs are available in the Connection Manager to set advanced connection settings. To establish a connection, you do not need to make modifications to these tabs; a connection can be established with the default settings. For more information on these advanced settings, see the chapter that corresponds with the tab name. For example, for information on the Verify tab, see [Verification](#) on page 9-1.

7. Click **Connect** to establish the connection.

NOTE: If the connection does not complete within 30 seconds, it is automatically canceled. If this timeout is not long enough for your environment, you can increase it by adjusting the **Communication Timeout** on the Configuration tab of the Management Console properties. Select **File, Options**, from the Management Console to access this screen.

8. Repeat steps 1 through 7 to connect multiple replication sets to the same target or to other targets, or to connect the same replication set to the same target or other targets.

Disconnecting a connection

To disconnect an established Storage Mirroring connection, right-click the connection on the right pane of the Management Console and select **Disconnect**. The source and target will be disconnected and in order to reconnect them, you must use the Connection Manager or Connection Wizard.

NOTE: If a connection is disconnected and the target is monitoring the source for failover, you will be prompted if you would like to continue monitoring for a failure. If you select **Yes**, the Storage Mirroring connection will be disconnected but the target will continue monitoring the source. To make modifications to the failure monitoring, you will need to use the Failover Control Center. See [Managing Failover Through the Failover Control Center](#) on page 11-17 for additional details. If you select **No**, the Storage Mirroring connection will be disconnected and the source will no longer be monitored for failure by the target.

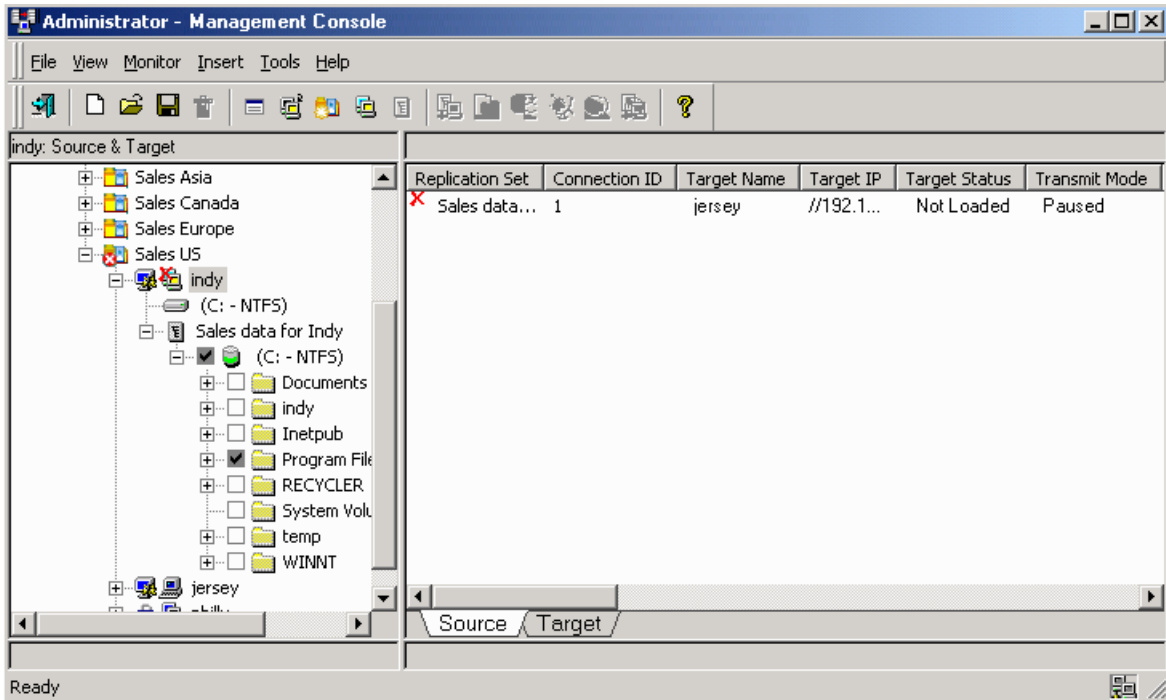
If a connection is disconnected while large amounts of data still remain in queue, the Management Console may become unresponsive while the data is being flushed. The Management Console will respond when all of the data has been flushed from the queue.

Monitoring a connection

You can monitor an established Storage Mirroring connection by viewing the connection statistics and/or the connection display. These are not the only methods available for monitoring an established Storage Mirroring connection. For information on all of the monitoring methods, see [Monitoring Tools](#) on page 13-1.

Connection statistics

- 1. To view the connection statistics on the Management Console, highlight the source machine on the left pane and the connection will appear on the right pane.
- 2. The right pane of the Management Console displays connection information. Use the horizontal scroll bar at the bottom of the right pane to view available connection statistics. For a complete list of the connection statistics, see [Management Console statistics](#) on page 13-2. For information on reordering the display of the connection statistics, see [Customizing the statistics](#) on page 13-4.





Connection display

You can monitor the connection by understanding the different icons displayed in the Management Console.

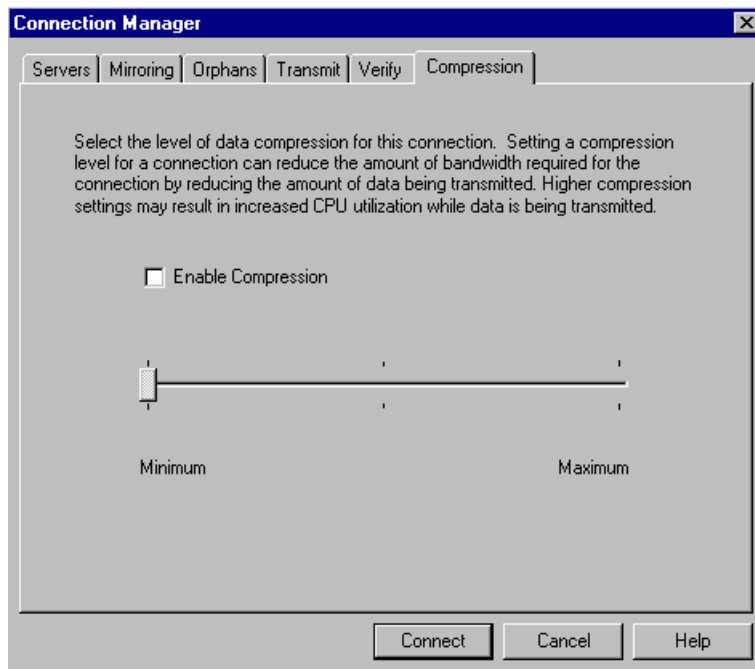
- ◆ **Established Connection Status Icon**—The source/target connection status appears as an icon to the left of the replication set name. (If you have reordered the connection data items displayed, the icon will appear to the left of the first column that is displayed.)
- ◆ **Established Connection Background Color Display**—The Storage Mirroring Management Console polls the source machine to determine the status of the source/target connection. The background color of an established connection identifies whether or not the Management Console and the source are communicating.

Connection Icons

Icons and Colors	Description
	A green checkmark appears in the right pane of the Management Console to the left of the connection data line for an established connection that is working properly.
	<p>A red X appears in the right pane of the Management Console to the left of a connection that contains an error. There are five cases that can lead to an error state:</p> <ul style="list-style-type: none">◆ The transmission has been broken◆ Replication is in a pending state◆ The target is not loaded◆ The target is busy◆ The source machine has reached the Stop Replicating Limit indicating that it is out of memory <p>To determine the exact problem, use the horizontal scroll bar at the bottom of the right pane of the Management Console and scroll to find the connection data item that appears in red. The connection problem is indicated in red.</p> <p>Note: The red X will only appear if replication is in a pending state due to processing. If replication is manually stopped, either by not initiating it when the connection is established or by manually stopping it, the connection will have a green check mark even though replication is not occurring.</p>
White	If the background is white, the Management Console has received a response from the source indicating that the source and Management Console are communicating.
Gray	If the background is gray, the number of missed responses from the source has been exceeded, therefore, communications between the Management Console and the source has been lost. The connection statistics are not updating when the background is gray. The background will return to white and the statistics will begin updating again when the Management Console receives a response from the source machine. For information on configuring the missed responses, see Customizing when the background color and icons change on page 13-9.

Configuring compression

1. You can configure compression when establishing a connection or after a connection has been created.
 - ◆ If you have not yet established a connection, follow the procedures as described in [Establishing a new connection](#) on page 6-8.
 - ◆ If you want to configure compression on an established connection, right-click the connection on the right pane of the Management Console and select **Connection Manager**.
2. Select the Compression tab to configure the compression settings.



3. By default, compression is disabled. To enable it, select **Enable Compression**.

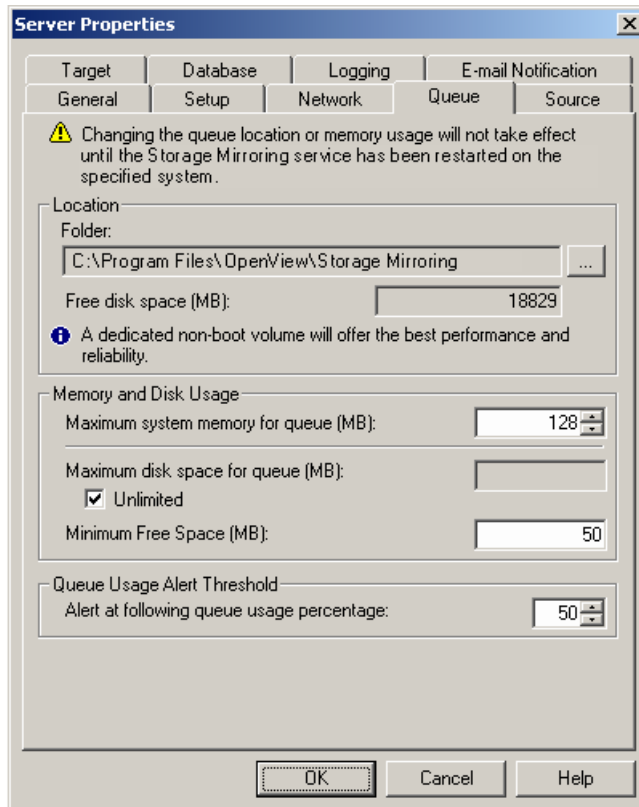
NOTE: You can specify compression for different source/target connections, but all connections to the same target will have the same compression settings.

4. Set the level from minimum to maximum compression.
5. When you are finished, click **Connect** or **OK**. (The button name will vary depending on whether the settings are for a new or existing connection.)

Configuring queuing

There are five queuing parameters available for configuration.

1. Right-click a source machine and select **Properties**.
2. Select the **Queue** tab.



3. Specify the **Folder** which will store the queue. For best results and reliability, you should select a dedicated, non-boot volume. The queue should be stored on a fixed, local NTFS volume.
4. Specify the **Maximum system memory for queue**. This is the amount of Windows system memory, in MB, that, when exceeded, will trigger queuing to disk. This value is dependent on the amount of physical memory available but has a minimum of 32 MB. By default, 128 MB of memory will be used.
5. Specify the **Maximum disk space for queue**. This is the maximum amount of disk space, in MB, in the specified **Folder** that can be used for Storage Mirroring queuing or you can select **Unlimited** which will allow the queue usage to automatically expand whenever the available disk space expands. When the disk space limit is reached, Storage Mirroring will automatically begin the auto-disconnect process. By default, Storage Mirroring will use the maximum amount of disk space available. Setting this value to zero (0) disables disk queuing.
6. Specify the **Minimum Free Space**. This is the minimum amount of disk space in the specified **Folder** that must be available at all times. By default, 50 MB of disk space will always remain free. The **Minimum Free Space** should be less than the amount of physical disk space minus **Maximum disk space for queue**.

NOTE: The **Maximum disk space for queue** and **Minimum Free Space** settings work in conjunction with each other. For example, assume your queues are stored on a 10 GB disk with the **Maximum disk space for queue** set to 10 GB and the **Minimum Free Space** set to 500 MB. If another program uses 5 GB, Storage Mirroring will only be able to use 4.5 GB so that 500 MB remains free.

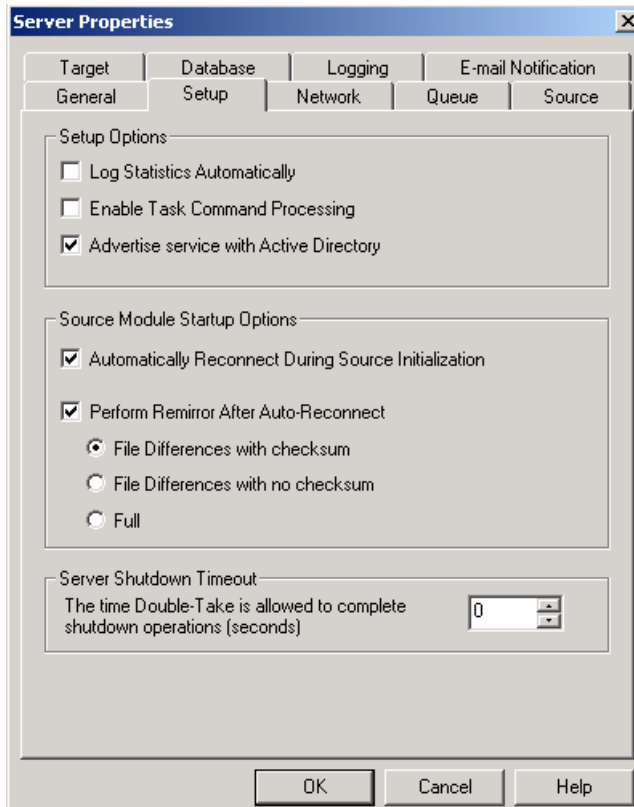
-
7. Specify a percentage for **Alert when queue usage reaches** to identify the amount of the queue that must be in use to trigger an alert message in the Windows Event Viewer. By default, the alert will be generated when the queue reaches 50%.
 8. Click **OK** to save the settings.

NOTE: Scanning the Storage Mirroring queue files for viruses can cause unexpected results. If anti-virus software detects a virus in a queue file and deletes or moves it, data integrity on the target cannot be guaranteed. As long as you have your anti-virus software configured to detect the actual production data, the anti-virus software can clean, delete, or move an infected file and the clean, delete, or move will be replicated to the target. This will keep the target from becoming infected and will not impact the Storage Mirroring queues.

Configuring auto-reconnect

If Storage Mirroring is stopped on a source machine while a Storage Mirroring connection is active, auto-reconnect will automatically reestablish the connection, without manual intervention by the administrator, when the Storage Mirroring source is available. When an auto-reconnect is performed, you may want to perform an auto-remirror to guarantee data integrity. Use the instructions below for both the auto-reconnect and auto-remirror options.

1. Right-click a source machine and select **Properties**.
2. Select the **Setup** tab.



3. Verify that the check box **Automatically Reconnect During Source Initialization** is marked to enable the auto-reconnect feature.
4. Verify that the check box **Perform Remirror after Auto-Reconnect** is marked to automatically remirror after an auto-reconnect. To understand how these options work together, see the chart [File Differences Mirror Options](#) on page 7-1.
5. Select the type of remirror you wish to perform.
6. Click **OK** to save these settings.

Managing Connections Through the Text Client

For each of the connection related commands, first identify which source you are working with by using the `source` command.

Command	<code>SOURCE</code>
Description	Identifies a machine as the active source machine
Syntax	<code><u>SOURCE</u> <source_machine></code>
Options	<code>source_machine</code> —Name of the machine
Examples	<code>source indy</code>

Establishing a new connection

1. If you do not know which replication set you will be connecting to the target machine, the `repset list` command will display the available replication sets for that source.

Command	<code>REPSET LIST</code>
Description	Lists all replication set names for the currently selected source
Syntax	<code>REPSET LIST</code>

```
Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)

- List of rep sets -      * DATABASE HAS BEEN MODIFIED *
exchange                  disabled
  (press any key)

> repset list
```

2. Connect the replication set to the target by using the `connect` command.

Command	<code>CONNECT</code>
Description	Establishes a connection between a replication set and a target machine
Syntax	<code>CONNECT <repset> TO <target_machine></code> <code>MAP EXACT</code> <code> MAP BASE <target_path></code> <code> MAP <source_path> TO <target_path> [, ...]</code> <code>[MIRROR NOMIRROR] [, REPLICATE NOREPLICATE] [, MONITOR NOMONITOR]</code> <code>[, ORPHANS NOORPHANS] [, COMPRESSION <level>] [CLEARRESTOREREQUIRED]</code>

Options

- ◆ **repset**—Name of the replication set
- ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address
- ◆ **MAP EXACT**—Specifies that the replication set data will be sent to the same logical volume on the target (c:\data and d:\files is copied to c:\data and d:\files, respectively)
- ◆ **MAP BASE target_path**—Substitute a complete path, including the volume, for **target_path** and the data will be replicated to target_path\SrcVolName on the target machine
- ◆ **MAP source_path TO target_path**—Custom location that specifies each directory on the source and where that data will be copied to on the target machine
- ◆ **...**—Indicates that the **source_path TO target_path** option can be used more than once for each source directory in the replication set
- ◆ **MIRROR**—Automatically initiates a mirror when the connection is established
- ◆ **NOMIRROR**—Does not initiate a mirror when the connection is established
- ◆ **REPLICATE**—Automatically initiates replication when the connection is established
- ◆ **NOREPLICATE**—Does not initiate replication when the connection is established
- ◆ **MONITOR**—Specifies that the target is going to monitor the specified source machine for failover. The source machine must have already been defined as a monitor machine.
- ◆ **NOMONITOR**—Specifies that the target is not going to monitor the source machine for failover
- ◆ **ORPHANS**—Removes orphan files on the target
- ◆ **NOORPHANS**—Does not remove orphan files on the target
- ◆ **COMPRESSION level**—Sets the compression level for data being sent to the target using levels 1 through 3 with 1 being minimum compression and 3 being maximum compression. This option sets the level only. It does not enable compression.
- ◆ **CLEARRESTOREREQUIRED**—Clears the restore required flag and initiates the connection

Examples

- ◆ connect Exchange to jersey map exact
- ◆ connect sql to jersey map base d:\DTFiles\

Notes

- ◆ The default settings for this command are mirror, replicate, nomonitor, and noorphans.
- ◆ If a path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.
- ◆ If you are establishing a connection within a NAT environment, you will need to specify the port of the router after the IP address (separated by a colon). See [NAT Environments](#) on page D-1 for more details.

Connecting to the Throughput Diagnostics Utility

To initiate a connection to the Throughput Diagnostics Utility, use the `connect to tdu` command. This logs the connection statistics to the file specified. The remaining connection flags are identical to the standard `connect` command.

Command	<code>CONNECT TDU</code>
Description	Establishes a connection between a replication set and the <u>T</u> hroughput <u>D</u> iagnostics <u>U</u> tility to imitate a normal connection without transmitting any data across the network
Syntax	<code>CONNECT <repset> TO TDU <file_name> [<i>connection flags</i>]</code>
Options	<ul style="list-style-type: none">◆ <i>repset</i>—Name of the replication set◆ <i>file_name</i>—Name of the file to store the connection statistics generated by the TDU◆ <i>connection flags</i>—The same options available in the standard <code>connect</code> command
Examples	<ul style="list-style-type: none">◆ <code>connect "Exchange Repset" to TDU</code>◆ <code>connect sql to TDU map c:\sql\data to e:\backup\sql\data</code>
Notes	The statistic file that the TDU creates can be viewed using DTStat. By default, the file is called <code>statistic.sts</code> . To view the statistic file, type <code>DTStat -f <filename></code> . For detailed information on DTStat, see DTStat on page 13-30.

Disconnecting a connection

Identify the connection ID by locating the number in the left column below the target machine name and use the `disconnect` command to disconnect the source/target connection.

Command	<code>DISCONNECT</code>
Description	Disconnects a specified source/target connection for the currently selected source.
Syntax	<code><u>DISCONNECT</u> <con_id *></code>
Options	<ul style="list-style-type: none">◆ <code>con_id</code>—Connection ID assigned to the source/target connection◆ <code>*</code>—Specifies all connection IDs.
Examples	<ul style="list-style-type: none">◆ <code>disconnect 1</code>◆ <code>disconnect *</code>

Monitoring a connection

Storage Mirroring offers a variety of methods for monitoring a Storage Mirroring connection. Statistics include information like connection up-time, target status, bytes queued or sent, etc. The informational data includes processing notifications, warnings or errors. Below is a brief description of how four of the monitoring methods are applicable to the Storage Mirroring connection. For detailed information on these and other monitoring methods, see [Monitoring Tools](#) on page 13-1.

- ◆ **Text Client**—The Storage Mirroring Text Client window displays a minimal number of statistical and informational data on screen.
- ◆ The DTCL `status` command displays statistical and informational data. Identify the connection ID by locating the number in the left column of the Text Client window and use the `status` command to display the information.

Command	<code>STATUS</code>
Description	Requests connection and statistical information
Syntax	<code>STATUS <u>CONNECT</u> MIRROR <u>REPLICATE</u> <u>TRANSMIT</u> <con_id></code>
Options	<ul style="list-style-type: none">◆ <code>CONNECT</code>—Displays connection information◆ <code>MIRROR</code>—Displays mirroring information◆ <code>REPLICATE</code>—Displays replication information◆ <code>TRANSMIT</code>—Displays transmission information◆ <code>con_id</code>—Connection ID assigned to the source/target connection
Examples	<ul style="list-style-type: none">◆ <code>status connect 1</code>◆ <code>status replicate 1</code>

- ◆ **LogViewer**—LogViewer is a utility within Storage Mirroring that gathers Storage Mirroring alerts and information messages. Alerts are processing notifications, warnings, and error messages. For detailed steps on how to use LogViewer, see [LogViewer](#) on page 13-39.
- ◆ **DTStat**—DTStats is a utility that takes snapshots of Storage Mirroring statistical data and writes the information to a file. For detailed steps on how to use DTStat, see [DTStat](#) on page 13-30.

Storing the connection ID in a variable

Variables can be used in DTCL commands to replace items that vary such as a machine name. A variable name must start with the \$ character and can contain letters, numbers, strings, and integers. Values are assigned to variables using the '=' assignment statement.

To store the connection ID assigned to an established connection in a variable, use the `conid` command.

Command	CONID
Description	Allows you to assign the value of a connection ID to a variable or list the target and replication set for all connection IDs for a source
Syntax	<ul style="list-style-type: none">◆ <code><variable>=CONID <repset> TO <target></code>◆ <code>CONID LIST <source_machine></code>
Options	<ul style="list-style-type: none">◆ variable—The name of the variable that you want to use to store the connection ID◆ repset—The replication set that was used to establish the connection◆ target—The target that was used to establish the connection◆ source_machine—The name of the source server
Examples	<pre>\$con_id=conid exchange to jersey \$ConnectionNumber=conid sql to indy conid list jersey</pre>
Notes	<ul style="list-style-type: none">◆ The <code>conid list</code> and <code>variable=conid</code> commands are two separate commands.◆ If no machine name is specified in the <code>conid list</code> command, the active source will be used.

For detailed information on other variables, see [Variables](#) on page A-45.

Configuring compression

1. To determine the current compression setting, use the `compression list` command. By default, compression is disabled.

Command	<code>COMPRESSION LIST</code>
Description	Identifies the compression level set and if compression is enabled
Syntax	<code>COMPRESSION LIST</code>

2. To set the level of compression, use the `compression set` command.

Command	<code>COMPRESSION SET</code>
Description	Sets the compression level
Syntax	<code>COMPRESSION SET <con_id> <level></code>
Options	<ul style="list-style-type: none">◆ <code>con_id</code>—Connection ID assigned to the source/target connection◆ <code>level</code>—Any whole number from 1 to 3 where 1 is minimum compression and 3 is maximum compression
Examples	<code>compression set 1 2</code>
Notes	◆ This command only sets the level of compression. It does not initiate compression.

3. To initiate compression, use the `compression enable` command.

Command	<code>COMPRESSION ENABLE</code>
Description	Enables compression
Syntax	<code>COMPRESSION ENABLE <con_id></code>
Options	<code>con_id</code> —Connection ID assigned to the source/target connection
Examples	<code>compression enable 1</code>

4. If you need to stop compression, use the `compression disable` command.

Command	<code>COMPRESSION DISABLE</code>
Description	Disables compression
Syntax	<code>COMPRESSION DISABLE <con_id></code>
Options	<code>con_id</code> —Connection ID assigned to the source/target connection
Examples	<code>compression disable 1</code>

Configuring queuing

There are five queuing parameters available for configuration. Use the `get` and `set` commands to view or modify an option. See [Storage Mirroring Commands](#) on page A-1 for complete command syntax.

- ◆ **QJournalDir**—This is the location where the queue is stored. For best results and reliability, you should select a dedicated, non-boot volume. The queue should be stored on a fixed, local NTFS volume.
- ◆ **QJournalFileSize**—This is the size, in MB, of each queuing transaction log file. By default, the file size is 5 MB.
- ◆ **QmemoryBufferMax**—This is the amount of Windows system memory, in MB, that, when exceeded, will trigger queuing to disk. This value is dependent on the amount of physical memory available but has a minimum of 32 MB. By default, 128 MB of memory will be used.
- ◆ **QJournalSpaceMax**—This is the maximum amount of disk space, in MB, in the specified QJournalDir that can be used for Storage Mirroring queuing. When this limit is reached, Storage Mirroring will automatically begin the auto-disconnect process. By default, Storage Mirroring will use an unlimited amount of disk space equivalent to 4,294,967,295 MB. This unlimited setting allows the disk queue usage to automatically expand whenever the available disk space expands. Setting this value to zero (0) disables disk queuing.
- ◆ **QJournalFreeSpaceMin**—This is the minimum amount of disk space, in MB, in the specified QJournalDir that must be available at all times. By default, 50 MB of disk space will always remain free. The QJournalFreeSpaceMin should be less than the amount of physical disk space minus QJournalSpaceMax.

NOTE: The **QJournalSpaceMax** and **QJournalFreeSpaceMin** settings work in conjunction with each other. For example, assume your queues are stored on a 10 GB disk with the **QJournalSpaceMax** set to 10 GB and the **QJournalFreeSpaceMin** set to 500 MB. If another program uses 5 GB, Storage Mirroring will only be able to use 4.5 GB so that 500 MB remains free.

- ◆ **QueueSizeAlertThreshold**—This is a percentage that identifies the amount of the queue that must be in use to trigger an alert message in the Windows Event Viewer. By default, the alert will be generated when the queue reaches 50%.

NOTE: Scanning the Storage Mirroring queue files for viruses can cause unexpected results. If anti-virus software detects a virus in a queue file and deletes or moves it, data integrity on the target cannot be guaranteed. As long as you have your anti-virus software configured to detect the actual production data, the anti-virus software can clean, delete, or move an infected file and the clean, delete, or move will be replicated to the target. This will keep the target from becoming infected and will not impact the Storage Mirroring queues.

Configuring auto-reconnect

If Storage Mirroring is stopped on a source machine while a Storage Mirroring connection is active, auto-reconnect will automatically reestablish the connection, without manual intervention by the administrator, when the Storage Mirroring source is available. When an auto-reconnect is performed, you may want to perform an auto-remirror to guarantee data integrity. Use the instructions below for both the auto-reconnect and auto-remirror options.

1. To determine the current auto-reconnect setting, use the `get` command with the `AutoReconnect` option.

Command	<code>GET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	<code>GET <option> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">◆ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.◆ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">◆ <code>get netport</code>◆ <code>get VerifyLogName</code>◆ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. To modify the auto-reconnect setting, use the `set` command with the `AutoReconnect` option. The `AutoReconnect` option is set to 1 for enabled or 0 for disabled. By default, auto-reconnect is enabled.

Command	<code>SET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	<code>SET <option>=<value> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">◆ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.◆ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.◆ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">◆ <code>set netport=1100</code>◆ <code>set VerifyLogName="HA server.log"</code>◆ <code>set DefaultAddress="129.12.34.3"</code>
Notes	<ul style="list-style-type: none">◆ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed.◆ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

3. To determine the current auto-remirror setting, use the `get` command with the `AutoRemirror` option.
4. To modify the Storage Mirroring auto-remirror setting, use the `set` command with the `AutoRemirror` option. The `AutoRemirror` option is set to 1 for enabled or 0 for disabled. By default, auto-remirror is enabled.

7 ► Mirroring

Mirroring is the process of transmitting the data contained in a replication set from the source to the target machine so that an identical copy of data exists on the target machine. All file and security attributes are also mirrored to the target machine. Mirroring must occur initially to generate a baseline copy from the source to the target. After mirroring has occurred, replication maintains an identical copy of the data on the target.

Mirroring Options

When configuring the mirroring process, you must specify what files you want sent from the source to the target using the following options:

- ◆ **Full Mirror**—A full mirror copies all of the directories and files in the replication set to the target machine. If a mirror has already been completed, another full mirror will overwrite the data on the target.
- ◆ **File Differences**—The file differences mirror option compares all files in the replication set on the source and their current state on the target machine. The comparison is based on the date, time, and size of the file.
 - ◆ **Only if newer on source**—If a file on the source is newer than the file on the target based on date and time stamps, the file is mirrored to the target machine. This option overwrites the file on the target machine. If the file on the target is newer than the file on the source, then the file is skipped.

NOTE: Accessing files on the target, directly or indirectly, could cause them to appear newer on the target. The Only if Newer on Source option will prevent newer files on the target from being successfully overwritten. For example, when a database application is started on the target, the date and time stamp on the file will be updated, although no data may actually change. If the Only if Newer on Source option is used, the database file on the target in this example, will not be updated, even though the source file may have newer data.

- ◆ **Block checksum**—This option compares the date, time, and size and for those files that are different, a checksum calculation comparison is performed. A checksum calculation is a formula applied to blocks of data to determine if the binary make-up of the block is identical. If the checksums on the source and target machine are the same, the block is skipped. If the checksums on the source and target machine are not the same, the block on the source is sent to the target. With this option, the entire file is not overwritten; only the block that is received from the source is overwritten.

NOTE: Mirroring does not find and delete files that exist on the target that are no longer on the source. See [Orphan Files](#) on page 7-3 for information on deleting files on the target that no longer exist in the source replication set.

The following table illustrates how the remirror options work together, as well as how they work with the global checksum setting on the Source tab of the Server Properties. Determine which options you need to select when configuring mirroring by identifying which type of action you want Storage Mirroring to perform during the mirror.

File Differences Mirror Options

Block Checksum All Server Properties Source Tab ^a	File Differences Mirroring Option	Only if Source is Newer Mirroring Option	Block Checksum Mirroring Option	Action Performed
(✓)	✓			Any file that is different on the source and target based on the date, time, and/or size is transmitted to the target. The mirror sends the entire file.
(✓)	✓	✓		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.

Block Checksum All Server Properties Source Tab ^a	File Differences Mirroring Option	Only if Source is Newer Mirroring Option	Block Checksum Mirroring Option	Action Performed
	✓		✓	Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
✓	✓		✓	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(✓)	✓	✓	✓	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

a. A checkmark enclosed in parentheses (✓) indicates that the global option on the Server Properties Source tab can be on or off. The use of this option does not change the action performed during the mirror.

NOTE: The file differences options used with auto-remirror vary slightly. See [Auto-Remirror](#) on page 7-4 for details.

Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the Block Checksum All option on the Server Properties Source tab to ensure proper file comparisons.

Stopping and Pausing Mirroring

While a mirror is in progress, you can manually stop or pause the transmission of data to the target. The Storage Mirroring connection will remain intact and the mirror data may or may not queue on the source depending on the option you choose.

- ◆ **Pause and Resume**—When pausing a mirror, Storage Mirroring stops queueing mirror data on the source but maintains a pointer to determine what information still needs to be mirrored to the target. Therefore, when resuming a paused mirror, the process continues where it left off.
- ◆ **Stop and Start**—When stopping a mirror, Storage Mirroring stops queueing mirror data on the source and does not maintain a pointer to determine what information still needs to be mirrored to the target. Therefore, when starting a mirror that has been stopped, the process will mirror all of the data contained in the replication set.

NOTE: Data integrity is not guaranteed until mirroring is complete.

Orphan Files

An orphan is a file that exists in the target location but is not in the source location. An orphan can be created when you delete a file contained in the source replication set while there is no Storage Mirroring connection. For example, if a connection was made and a mirror was completed and then the connection was stopped and a file was deleted on the source, an orphan file will exist on the target. Because the connection has been disconnected, the delete operation is not replicated to the target and the file is not deleted.

The Storage Mirroring orphan feature allows you to move or delete orphan files that appear on the target but not on the source. You can specify to move or delete all orphan files or to only those orphans that are older than a specified period of time.

When the orphans feature is enabled, the target is checked for orphan files when a mirror, verify, or restore is performed. You can also check for orphan files manually. The results of orphan processing are maintained in the LogViewer log files on the target including the number of moved/deleted orphan files, the directories, and the number of bytes.

NOTE: Since the orphans feature will (re)move all files that exist on the target location that are not on the source, if you manually copy files and/or directories into the target location that are not also in the source location, these files will be considered orphans and will be (re)moved. For example, suppose you have the following rules defined in a replication set:

```
d:\          include, recursive
d:\users     include, recursive
d:\data      include, recursive
d:\temp      exclude, recursive
```

Files or directories that exist on the target under `d:\`, `\users`, or `\data` that do not exist on the source will be (re)moved. Since `\temp` is specifically excluded, it is not considered an orphan.

The orphans feature does not move or delete alternate data streams. To do this, use a full mirror which will delete the additional stream(s) when the file is re-created.

Auto-Remirror

A remirror repeats the mirroring process to guarantee the integrity of the data on the target. You can manually remirror at any time you desire and you can also configure an automatic remirror, called auto-remirror, to occur after an auto-disconnect. For more information on auto-disconnect, see [Auto-Disconnect and Auto-Reconnect](#) on page 6-7 .

NOTE: Auto-remirror cannot be set on a connection basis. If auto-remirror is turned on, all connections from that source machine will perform an auto-remirror on a reconnect.

If auto-remirror is disabled and an auto-reconnect occurs, the transmission state of the connection will remain pending after the reconnect until a remirror is started manually.

The following table illustrates how the auto-remirror options work together, as well as how they work with the global checksum setting on the Source tab of the Server Properties. Determine which options you need to select when configuring auto-remirror by identifying which type of action you want Storage Mirroring to perform during the auto-remirror.

File Differences Auto-Remirror Options


Block Checksum All ^a	Checksum	No Checksum	Action Performed
(✓)		✓	Any file that is different on the source and target based on the date, time, and/or size is transmitted to the target. The mirror sends the entire file.
	✓		Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
✓	✓		The mirror performs a checksum comparison on all files and only sends those blocks that are different.

a. This option is located on the Server Properties Source tab. A checkmark enclosed in parentheses (✓) indicates that the global option on the Server Properties source tab can be on or off. The use of this option does not change the action performed during the mirror.

NOTE: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the **Block Checksum All** option on the Server Properties Source tab to ensure proper file comparisons.

How Mirroring Works

To better understand how Storage Mirroring mirroring works, the following diagram identifies the different steps that are completed when a mirror is performed. Each numbered step is described below the diagram.



1. Mirroring is initiated by the user, either manually through one of the clients or automatically when the connection is created.
2. Storage Mirroring determines which data needs to be sent to the target depending on the mirroring criteria that was specified through the client. If it is a full mirror, all of the files are sent to the target. If it is a file differences mirror, the files contained in the replication set on the source are compared against the copy of the replication set on the target to determine which files need to be mirrored.
3. Storage Mirroring transmits the mirror data to the target machine.
4. As each packet of mirror data is received on the target, the target returns an acknowledgment to the source confirming that the mirrored data has been received.

Managing Mirroring Using the Management Console

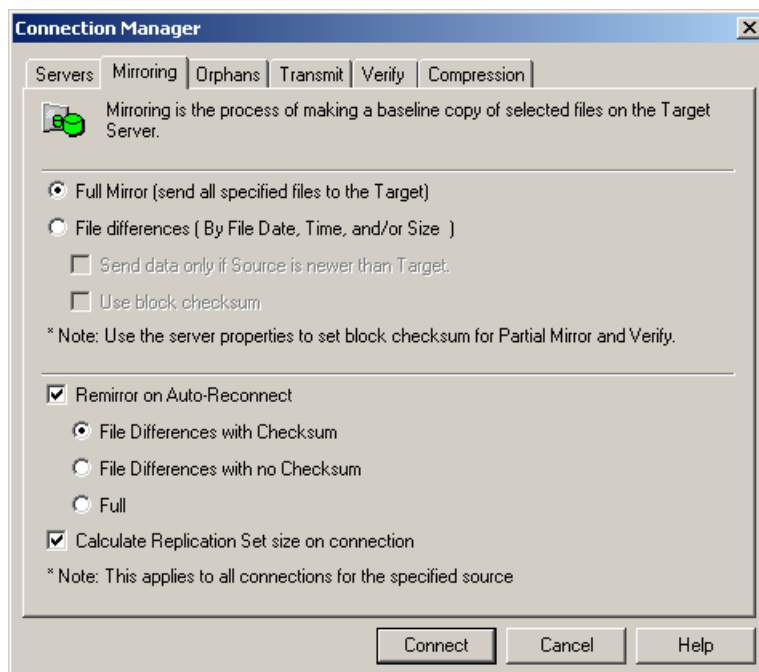
Through the Management Console, a mirror can be initiated upon connection, manually controlled, and monitored. In addition, the Management Console can be used to configure orphan files and to automatically run an auto-remirror after an auto-reconnect to ensure that data is identical on the source and target.

Initiating a mirror upon connection

1. There are four methods available for establishing a connection and opening the Connection Manager:
 - ◆ Highlight a replication set and select **Tools, Connection Manager**.
 - ◆ Right-click on a replication set and select **Connection Manager**.
 - ◆ Drag and drop a replication set to a target machine on the left pane of the Management Console.
 - ◆ Highlight a replication set and then drag and drop the replication set to a target machine on the right pane of the Management Console.

The Connection Manager will open and the **Servers** tab will be selected. Complete the **Servers** tab entries as described in [Establishing a new connection](#) on page 6-8.

2. Select the **Mirroring** tab.

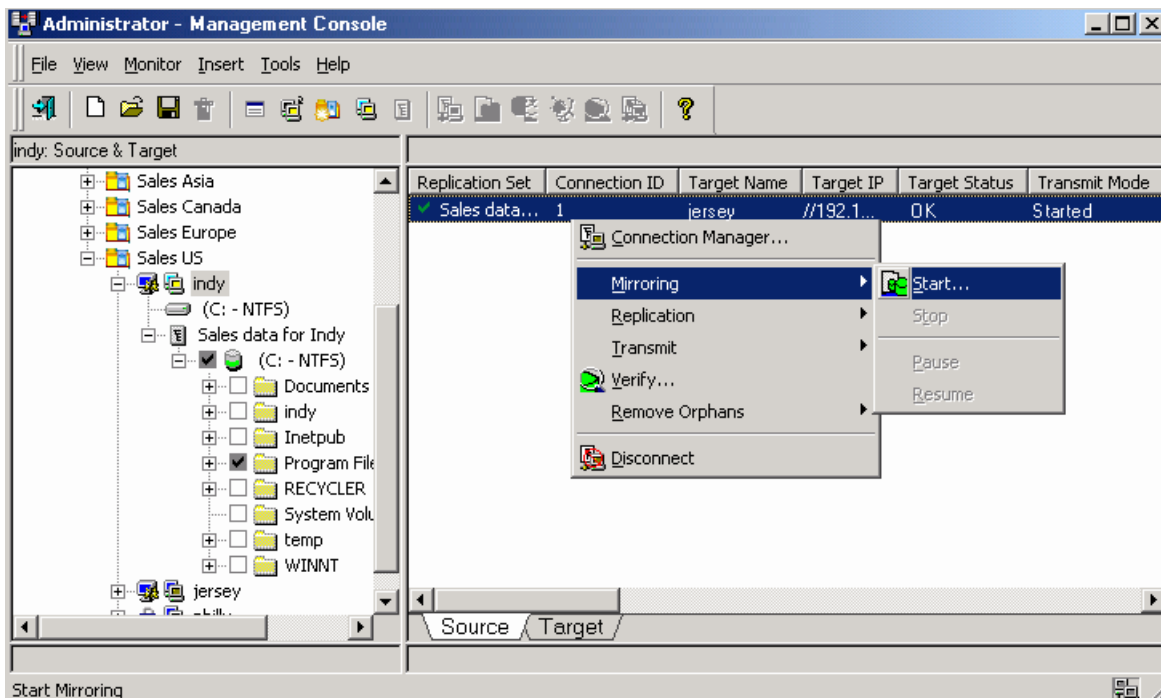


3. Select the type of mirror, either **Full** or **File Differences**, to perform.
4. If you select **File differences**, you may additionally select one or both of the following options: **Send data only if Source is newer than Target** and/or **Use block checksum**. To understand how these options work together, see the chart [File Differences Mirror Options](#) on page 7-1.
5. Enable or disable the auto-remirror on auto-reconnect feature by marking or clearing the **Remirror on Auto-Reconnect** check box. Specify the type of auto-remirror that will be performed. To understand how these options work together, see the chart [File Differences Mirror Options](#) on page 7-1.
6. By default, the **Calculate Replication Set size on connection** option is enabled. To disable this feature, clear the check box. For detailed information on this feature, see [Calculating replication set size](#) on page 5-10.
7. Click **Connect** when you are finished.


NOTE: Other tabs are available in the Connection Manager to set advanced connection settings. To establish a connection with mirroring options, you do not need to make modifications to these tabs; a connection can be established with the default settings. For more information on these advanced settings, see the chapter that corresponds with the tab name. For example, for information on the Verify tab, see [Verification](#) on page 9-1.

Controlling a mirror manually

After a connection is established you can manually start, stop, pause, or resume mirroring. Right-click the connection on the right pane of the Management Console and select **Mirroring** and the appropriate mirror control.



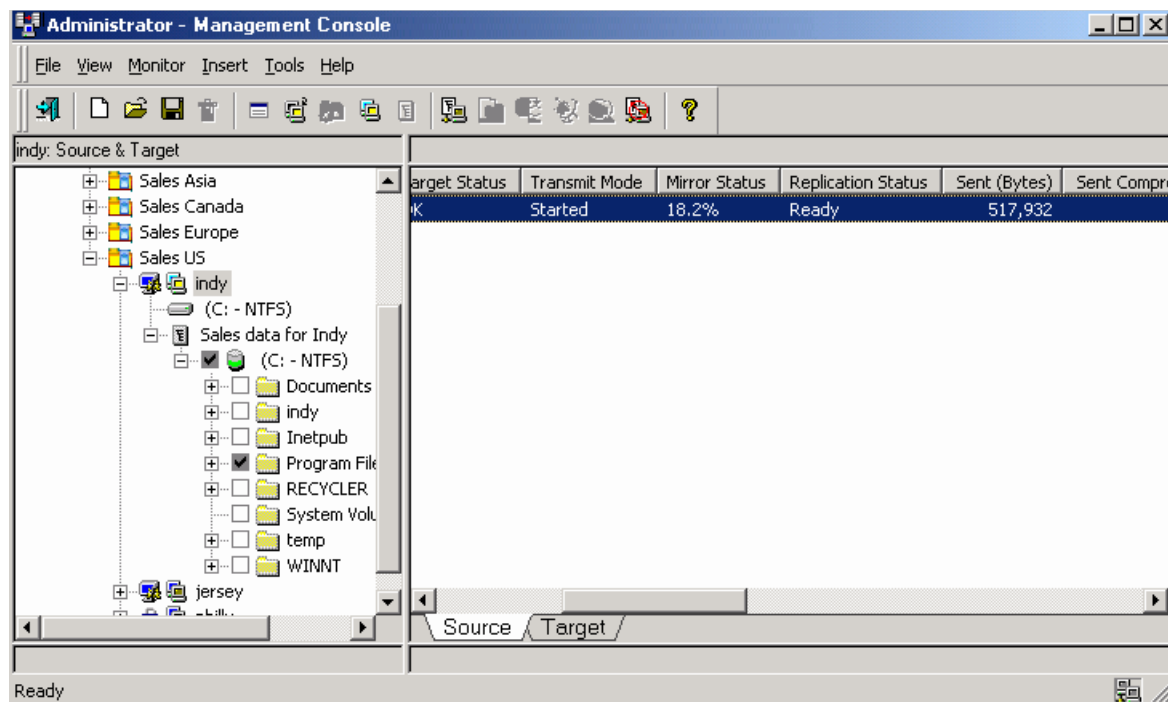
- ◆ **Start**—If you select to start a mirror, you will need to make the following two selections on the Start Mirror dialog box:



- ◆ **Type of Mirror**—Specify the type of mirror, **Full** or **File Differences**, you wish to perform. To understand how the File Differences options work together, see the chart [File Differences Mirror Options](#) on page 7-1.
- ◆ **Replication Set Size Calculation**—By default, the **Calculate Replication Set size prior to mirror** option is enabled. To disable this feature, clear the check box. For detailed information on this feature, see [Calculating replication set size](#) on page 5-14.
- ◆ **Stop, Pause or Resume**—These selections will automatically initiate the selected mirror control.

Monitoring a mirror

1. To view the connection on the Management Console, highlight the source machine on the left pane and the connection will appear on the right pane. Use the horizontal scroll bar at the bottom of the right pane to view the **Mirror Status** column; the status of the mirroring operation for this connection is displayed. The possible states are:
 - ◆ **Idle**—No mirroring activity is taking place.
 - ◆ **Mirroring**—If the size of the replication set was not calculated prior to the connection being established, the status indicates **Mirroring**.
 - ◆ **Percentage Complete**—If the size of the replication set was calculated prior to the connection being established, the percentage of the mirror that is complete will be displayed.
 - ◆ **Paused**—No new mirror data is being written to the mirror queue.
 - ◆ **Waiting**—The source has finished sending the data to the target, but the target has not finished writing to the disk.
 - ◆ **Verifying**—The source is verifying if the data on the source and target are identical.
 - ◆ **Removing Orphans**—Storage Mirroring is checking for orphan files within the target path location that exist on the target but not on the source. These files will be removed.



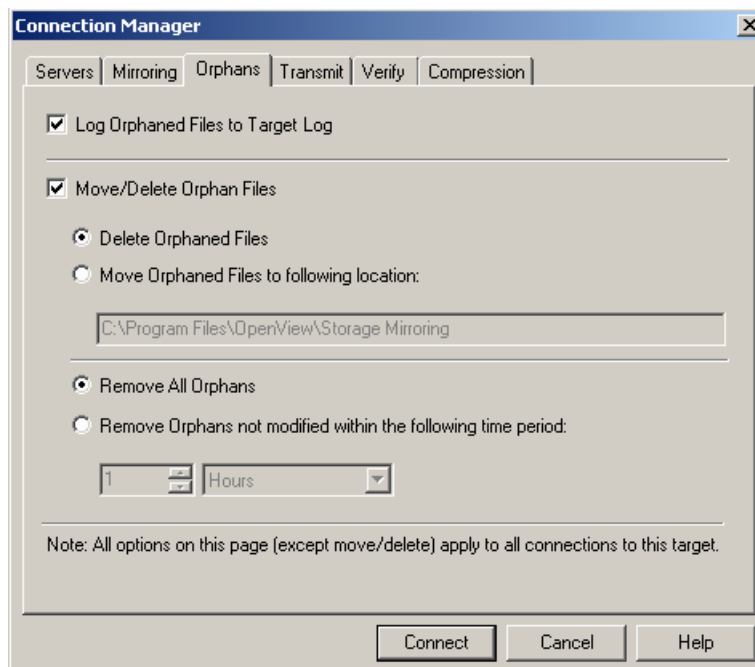
2. Additional mirroring statistics are also available by using the horizontal scroll bar to display the following columns:
 - ◆ **Sent (Bytes)**—The total number of mirror and replication bytes that have been sent during this connection.
 - ◆ **Sent Mirror (Bytes)**—The total number of mirror bytes only that have been sent during this connection.
 - ◆ **Skipped Mirror (Bytes)**—The total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is the same on the source and target machines.
 - ◆ **Remaining Mirror (Bytes)**—The total number of mirror bytes only that remain to be sent to the target.

NOTE: There are other statistics available through the Management Console, but they are not displayed by default. To change the statistics display, see [Customizing the statistics](#) on page 13-4.

You may see replication operations during a mirror, even if you are not currently replicating any data. These operations are pre-allocating space in order to avoid fragmentation on the target. These operations are normal.

Configuring orphan files

1. You can configure orphan files when establishing a connection or after a connection has been created.
 - ◆ If you have not yet established a connection, follow the procedures as described in [Establishing a new connection](#) on page 6-8.
 - ◆ If you want to configure the orphan files setting on an established connection, right-click the connection on the right pane of the Management Console and select **Connection Manager**.
 - ◆ If you want to configure orphan files during a restoration, select **Tools, Restoration Manager**. The Restoration Manager has the same Orphans tab as shown in the Connection Manager.
2. Select the Orphans tab to configure the orphan files feature.



3. By default, the orphan files feature is disabled. To enable it, mark **Move/Delete Orphan Files**.

NOTE: You can specify to enable or disable orphans for different source/target connections, but all connections to the same target will have the same orphan settings.

4. Specify if you want to log the name of the orphan files to the Storage Mirroring log file on the target by marking **Log Orphaned Files to Target Log**.
5. Specify if you want to **Delete Orphaned Files** or **Move Orphaned Files** to a different location. If you select the move option, identify the location where these orphan files will be located.

NOTE: If you are moving files, make sure the directory you specify to move the files to is not included in the destination of the replication set data so that the orphan files are only moved once.

6. Specify if you want to **Remove All Orphans** or **Remove Orphans not modified within the following time period**. If you select the time-based option, only orphans older than the time you specify will be removed.
7. When you are finished, click **Connect** or **OK**. (The button name will vary depending on whether the settings are for a new or existing connection.)

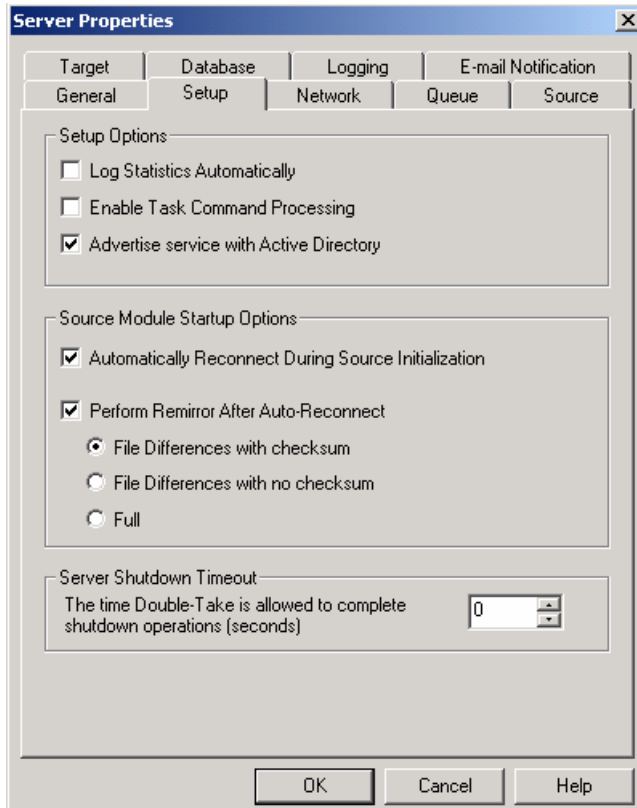
NOTE: If you want to remove orphan files manually, right-click an established connection and select **Remove Orphans, Start**. If you want to stop the process after it has been started, right-click the connection and select **Remove Orphans, Stop**.

Setting auto-remirror

In the event of a source machine failure or if the Storage Mirroring service stops, data integrity between the source and target cannot be guaranteed. To resolve this issue, the Storage Mirroring auto-reconnect process automatically reestablishes any source/target connections that existed prior to Storage Mirroring being unavailable. After an auto-reconnect, an auto-remirror will ensure that data is identical on the source and target machines.

You can establish the auto-remirror settings when establishing the initial connection or you can establish it through the Server Properties.

1. Right-click a machine name in the left pane of the Management Console and select **Properties**.
2. Select the **Setup** tab.



3. Verify that the **Perform Remirror After Auto-Reconnect** check box is selected to initiate auto-remirror after an auto-reconnect. Specify the type of mirror that you wish to perform. To understand how these options work together, see the chart [File Differences Mirror Options](#) on page 7-1.

NOTE: Auto-remirror cannot be set on a connection basis. If auto-remirror is turned on, all connections from that source machine will perform an auto-remirror on a reconnect.

Managing Mirroring Using the Text Client

For each of the mirroring commands, first identify which source you are working with by using the `source` command.

Command	<code>SOURCE</code>
Description	Identifies a machine as the active source machine
Syntax	<code>SOURCE <source_machine></code>
Options	<code>source_machine</code> —Name of the machine
Examples	<code>source indy</code>

Initiating a mirror on connection

1. If you do not know which replication set you will be connecting to the target machine, the `repset list` command will display the available replication sets for that source.

Command	<code>REPSET LIST</code>
Description	Lists all replication set names for the currently selected source
Syntax	<code>REPSET LIST</code>

2. Connect the replication set to the target by using the `connect` command. To initiate mirroring when the connection is established, specify the `mirror` switch with this command. By default, mirroring is started automatically.

Command	<code>CONNECT</code>
Description	Establishes a connection between a replication set and a target machine
Syntax	<pre>CONNECT <repset> TO <target_machine> MAP EXACT MAP BASE <target_path> MAP <source_path> TO <target_path> [,...] [MIRROR NOMIRROR] [, REPLICATE NOREPLICATE] [, MONITOR NOMONITOR] [, ORPHANS NOORPHANS] [, COMPRESSION <level>] [CLEARRESTOREREQUIRED]</pre>

Options

- ◆ **repset**—Name of the replication set
- ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address
- ◆ **MAP EXACT**—Specifies that the replication set data will be sent to the same logical volume on the target (c:\data and d:\files is copied to c:\data and d:\files, respectively)
- ◆ **MAP BASE target_path**—Substitute a complete path, including the volume, for **target_path** and the data will be replicated to target_path\SrcVolName on the target machine
- ◆ **MAP source_path TO target_path**—Custom location that specifies each directory on the source and where that data will be copied to on the target machine
- ◆ **...**—Indicates that the **source_path TO target_path** option can be used more than once for each source directory in the replication set
- ◆ **MIRROR**—Automatically initiates a mirror when the connection is established
- ◆ **NOMIRROR**—Does not initiate a mirror when the connection is established
- ◆ **REPLICATE**—Automatically initiates replication when the connection is established
- ◆ **NOREPLICATE**—Does not initiate replication when the connection is established
- ◆ **MONITOR**—Specifies that the target is going to monitor the specified source machine for failover. The source machine must have already been defined as a monitor machine.
- ◆ **NOMONITOR**—Specifies that the target is not going to monitor the source machine for failover
- ◆ **ORPHANS**—Removes orphan files on the target
- ◆ **NOORPHANS**—Does not remove orphan files on the target
- ◆ **COMPRESSION level**—Sets the compression level for data being sent to the target using levels 1 through 3 with 1 being minimum compression and 3 being maximum compression. This option sets the level only. It does not enable compression.
- ◆ **CLEARRESTOREREQUIRED**—Clears the restore required flag and initiates the connection

Examples

- ◆ connect Exchange to jersey map exact
- ◆ connect sql to jersey map base d:\DTFiles\

Notes

- ◆ The default settings for this command are mirror, replicate, nomonitor, and noorphans.
- ◆ If a path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.
- ◆ If you are establishing a connection within a NAT environment, you will need to specify the port of the router after the IP address (seperated by a colon). See [NAT Environments](#) on page D-1 for more details.

Controlling a mirror manually

Identify the connection ID by locating the number in the left column below the target machine name and use one of the following commands for the appropriate mirror control:

◆ Mirror Stop

Command	MIRROR STOP
Description	Stops a paused mirror
Syntax	<u>MIRROR</u> STOP <con_id *>
Options	<ul style="list-style-type: none">◆ con_id—Connection ID assigned to the source/target connection◆ *—Specifies all connection IDs
Examples	<ul style="list-style-type: none">◆ <code>mirror stop 1</code>◆ <code>mirror stop *</code>

◆ Mirror Start

Command	MIRROR START
Description	Initiates the mirror process
Syntax	<u>MIRROR</u> START <con_id> [CLEAN DIFFERENT [,NEWER] ,CHECKSUM NOCHECKSUM] [ORPHANS NOORPHANS] [<u>CALCULATE</u> NOCALCULATE] [CLEARRESTOREREQUIRED]
Options	<ul style="list-style-type: none">◆ con_id—Connection ID assigned to the source/target connection◆ CLEAN—Deletes all files in the specified target path prior to mirroring◆ DIFFERENT—Mirrors only those files that are different based on the file date, time, and/or size◆ NEWER—Mirrors only those files that are newer on the source than on the target◆ CHECKSUM—Mirrors only those blocks that are different based on block checksum comparisons◆ NOCHECKSUM—Does not perform a checksum comparison when mirroring files◆ ORPHANS—Removes orphan files on the target◆ NOORPHANS—Does not remove orphan files on the target◆ CALCULATE—Calculate the size of the replication set prior to mirroring◆ CLEARRESTOREREQUIRED—Clears the restore required flag and initiates the mirror
Examples	<ul style="list-style-type: none">◆ <code>mirror start 1 different, newer</code>◆ <code>mirror start 3 clean</code>◆ <code>mirror start 2 different, checksum</code>
Notes	<ul style="list-style-type: none">◆ In order to use the <code>clean</code> option, you must have the cleaner program setting enabled. For detailed information, see the cleaner option in the table Storage Mirroring Text Client program settings on page 14-15.◆ By default, orphan files will not be removed and the size of the replication set database will be calculated.

◆ Mirror Pause

Command `MIRROR PAUSE`

Description Pauses a mirror that is in progress

Syntax `MIRROR PAUSE <con_id | *>`

Options ◆ `con_id`—Connection ID assigned to the source/target connection
 ◆ `*`—Specifies all connection IDs.

Examples ◆ `mirror pause 1`
 ◆ `mirror pause *`

◆ Mirror Resume

Command `MIRROR RESUME`

Description Resumes a paused mirror

Syntax `MIRROR RESUME <con_id | *>`

Options ◆ `con_id`—Connection ID assigned to the source/target connection
 ◆ `*`—Specifies all connection IDs.

Examples ◆ `mirror resume 1`
 ◆ `mirror resume *`

Monitoring a mirror

Storage Mirroring offers a variety of methods for monitoring Storage Mirroring mirroring. Statistics include information like connection up-time, target status, bytes queued or sent, and so on. The informational data includes processing notifications, warnings or errors. Below is a brief description of how four of the monitoring methods are applicable to the Storage Mirroring connection. For detailed information on these and other monitoring methods, see [Monitoring Tools](#) on page 13-1.

- ◆ **Text Client**—The Storage Mirroring Text Client window displays a minimal number of statistical and informational data on screen.
- ◆ The DTCL `status` command displays statistical and informational data. Identify the connection ID by locating the number in the left column of the Text Client window and use the `status` command to display the information.

Command	<code>STATUS</code>
Description	Requests connection and statistical information
Syntax	<code>STATUS <u>CONNECT</u> MIRROR <u>REPLICATE</u> <u>TRANSMIT</u> <con_id></code>
Options	<ul style="list-style-type: none">◆ <code>CONNECT</code>—Displays connection information◆ <code>MIRROR</code>—Displays mirroring information◆ <code>REPLICATE</code>—Displays replication information◆ <code>TRANSMIT</code>—Displays transmission information◆ <code>con_id</code>—Connection ID assigned to the source/target connection
Examples	<ul style="list-style-type: none">◆ <code>status connect 1</code>◆ <code>status replicate 1</code>

- ◆ **LogViewer**—LogViewer is a utility within Storage Mirroring that gathers Storage Mirroring alerts and information messages. Alerts are processing notifications, warnings, and error messages. For detailed steps on how to use LogViewer, see [LogViewer](#) on page 13-39.
- ◆ **DTStat**—DTStats is a utility that takes snapshots of Storage Mirroring statistical data and writes the information to a file. For detailed steps on how to use DTStat, see [DTStat](#) on page 13-30.

NOTE: You may see replication operations during a mirror, even if you are not currently replicating any data. These operations are pre-allocating space in order to avoid fragmentation on the target. These operations are normal.

Configuring orphan files

1. To determine the current orphan files setting, use the `get` command with the `MoveOrphanedFiles` option. By default, orphan files are not deleted.

Command	GET
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	GET <option> [<i>machine</i>]
Options	♦ <i>option</i> —See Storage Mirroring Text Client program settings on page 14-15. ♦ <i>machine</i> —Name of the machine
Examples	♦ <code>get netport</code> ♦ <code>get VerifyLogName</code> ♦ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. To modify the orphan files setting, use the `set` command with the `MoveOrphanedFiles` option. The `MoveOrphanedFiles` option is set to 0 to delete orphan files or 1 to move orphan files.

Command	SET
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	SET <option>=<value> [<i>machine</i>]
Options	♦ <i>option</i> —See Storage Mirroring Text Client program settings on page 14-15. ♦ <i>value</i> —See Storage Mirroring Text Client program settings on page 14-15. ♦ <i>machine</i> —Name of the machine
Examples	♦ <code>set netport=1100</code> ♦ <code>set VerifyLogName="HA server.log"</code> ♦ <code>set DefaultAddress="129.12.34.3"</code>
Notes	♦ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed. ♦ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

3. You can also delete or move only those orphan files that are older than a specified time. To determine the current remove orphans time, use the `get` command with the `RemoveOrphansTime` option.
4. To modify the length of time applied to determine which orphan files are deleted or moved, use the `set` command with the `RemoveOrphansTime` option. This option accepts any valid number of minutes between 1 and 131072. By default, the length of time is 60 minutes.
5. If you are moving the orphan files, these files are relocated to another directory. To determine the location where the orphaned files are moved, use the `get` command with the `MoveOrphansDir` option.

-
6. To modify the location of renamed orphan files, use the `set` command with the `MoveOrphansDir` option. This option accepts any valid volume and directory. By default, the directory is the same drive where Storage Mirroring is running in a directory called `\RemoveOrphans`.

NOTE: If you are moving files, make sure the directory you specify to move the files to is not included in the destination of the replication set data so that the orphan files are only moved once.

7. To determine if orphaned files are being logged on the target, use the `get` command with the `LogAllOrphans` option.
8. To modify whether or not the orphaned files are logged to the Storage Mirroring log on the target, use the `set` command with the `LogAllOrphans` option. This option is set to 0 to not log the files or 1 to log the files.
9. Orphan files will be removed or renamed, depending on your settings, when a mirror, verify, or restore is performed. If you want to remove orphan files manually, use the `orphans start` command.

Command `ORPHANS START`

Description Manual operation to remove any orphan files on the target

Syntax `ORPHANS START con_id [CLEARRESTOREREQUIRED]`

Options ♦ *con_id*—Connection ID assigned to the source/target connection
 ♦ `CLEARRESTOREREQUIRED`—Clears the restore required flag and initiates the orphan operation

Examples `orphans start 1`

10. If you want to stop removing orphan files after it has been started, use the `orphans stop` command.

Command `ORPHANS STOP`

Description Stops the process to remove orphan files on the target

Syntax `ORPHANS STOP con_id`

Options *con_id*—Connection ID assigned to the source/target connection

Examples `orphans stop 1`

Setting auto-remirror

In the event of a source machine failure or if the Storage Mirroring service stops, data integrity between the source and target cannot be guaranteed. To resolve this issue, the Storage Mirroring auto-reconnect process automatically reestablishes any source/target connections that existed prior to Storage Mirroring being unavailable. After an auto-reconnect, an auto-remirror will ensure that data is identical on the source and target machines. The default installation will configure the source to perform an auto-remirror if needed.

1. To determine the current auto-remirror setting, use the `get` command with the `AutoRemirror` option.

Command	<code>GET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	<code>GET <option> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>get netport</code>♦ <code>get VerifyLogName</code>♦ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. To modify the auto-remirror setting, use the `set` command with the `AutoRemirror` option. The `AutoRemirror` option is set to 1 for enabled or 0 for disabled. By default, auto-remirror is enabled.

Command	<code>SET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	<code>SET <option>=<value> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>set netport=1100</code>♦ <code>set VerifyLogName="HA server.log"</code>♦ <code>set DefaultAddress="129.12.34.3"</code>
Notes	<ul style="list-style-type: none">♦ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed.♦ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

8 ► Replication

A file change, also called a file request or operation, is the creation, modification, or deletion of a file. This also includes file attribute and permission changes. Replication is the real-time transmission of these file changes as they occur. The file changes are replicated from the source to the target, maintaining an identical copy of data on the target machine.

Replication is the key to Storage Mirroring; without having identical copies of data on the source and target machines at all times, the target would not be able to meet the high availability and disaster recovery needs of your complex network environment.

Replication Options


Replication can be started automatically on connection, which is the default and recommended configuration, and it can also be started and stopped manually at any time. Stopping replication is separate from stopping the transmission of data to the target. When replication stops, data is not added to the queue on the source.

WARNING: When stopping replication, source/target data integrity is no longer guaranteed. A remirror will be necessary to guarantee that the source and target data are identical.

You can also turn off replication when you are not concerned about real-time backup, but want to use the Storage Mirroring block checksum mirror to periodically synchronize files.

How Replication Works

To better understand how Storage Mirroring replication works, the following diagram identifies the steps that occur when data is replicated from the source to the target machine.



1. The file system handles all file requests when an application creates, modifies, or deletes data on the source machine.
2. The file requests are filtered by the Storage Mirroring File System Driver.
3. The Storage Mirroring File System Driver allows all file requests to be forwarded to the file system and to Storage Mirroring.
4. The file system writes the operation to disk on the source machine.
5. Storage Mirroring converts the file requests into replication packets.
6. The Storage Mirroring source sends the replication packets to the Storage Mirroring target where they are applied to the target copy of the data.

Recovery From Network Losses

Storage Mirroring handles packet losses on the network as well. Storage Mirroring verifies the incoming sequencing before looking at execution ordering. Here is an example.

1. The source sends packets 1, 2, 3, and 4 to the target.
2. The target receives packet 1. It expects 1, so it passes the operation to the holding/execution handlers and sends its acknowledgement of packet 1 back to the source.
3. Packet 2 is lost on the network.
4. The target receives packet 3. It expects 2, so it discards packet 3 and resends its acknowledgement of operation 1 to the source because it was the last successful operation.
5. Because the source already received an acknowledgement for packet 1, it puts any operations that are in the acknowledgement queue (starting with operation 2) back on the front of the transmit queue to be sent again, in order.

Windows Permissions

If you are using a Windows domain security model by assigning users at the domain level, each user is assigned a security ID (SID). When Storage Mirroring replicates a file to the target, the SID is also replicated. If you are using a Windows local security model by assigning users at the local level, users that appear on multiple machines each have different SIDs. After a Storage Mirroring file and SID is replicated, a local user will not be able to access the file because the user's SID on the target machine is different from the SID that was replicated from the source machine.

Storage Mirroring allows you to replicate Windows permission attributes by local name as well as security ID (SID). By replicating Windows security by name, you can transmit the owner name with the file. If that user exists on the target, then the SID associated with the user will be applied to the target file ownership. If that user does not exist on the target, then the ownership will be unknown.

By default, **Replicate NT Security by Name** is disabled. If you are using a Windows local security model, you will need to enable this feature so that users can access the data on the target.

NOTE:	If you enable Replicate NT Security by Name , make sure that the same groups and users exist on the target as they do on the source. Otherwise, mirror and replication speed will be slower while adjustments are made for the missing groups or users.
--------------	--

Task Command Processing

Task command processing is a Storage Mirroring feature that allows you to insert and run tasks at various points during the replication of data. Because the tasks are user-defined, you can achieve a wide variety of goals with this feature. For example, you might insert a task to create a snapshot or backup on the target after a certain segment of data from the source has been applied on the target.

Task command processing can be enabled from both the Management Console and the text clients but can only be initiated using the DTCL `queuetask` command through the text clients.

Managing Replication Through the Management Console

Using the Management Console, replication can be initiated upon connection, manually controlled, and monitored. In addition, Windows permissions can be replicated by name and task commands can be inserted in the replication queue.

Initiating replication upon connection

1. There are four methods available for establishing a connection and opening the Connection Manager:
 - ◆ Highlight a replication set and select **Tools, Connection Manager**.
 - ◆ Right-click on a replication set and select **Connection Manager**.
 - ◆ Drag and drop a replication set to a target machine on the left pane of the Management Console.
 - ◆ Highlight a replication set and then drag and drop the replication set to a target machine on the right pane of the Management Console.

The Connection Manager will open and the **Servers** tab will be selected.

2. Complete the Servers tab entries as described in [Establishing a new connection](#) on page 6-8.

The screenshot shows the 'Connection Manager' dialog box with the 'Servers' tab selected. The dialog has several tabs: 'Servers', 'Mirroring', 'Orphans', 'Transmit', 'Verify', and 'Compression'. The 'Servers' tab contains instructions: 'Select Source Server, Replication Set, Target Server, and Route. Define the destination paths and connection options.' Below this, there are four dropdown menus: 'Source Server' (set to 'indy'), 'Replication Set' (set to 'Sales data for Indy'), 'Target Server' (set to 'jersey'), and 'Route' (set to '192.168.2.31 : 1100'). Below these is a 'Mappings' section with two radio buttons: 'All To One' (selected) and 'One To One'. Below the radio buttons is a table with two columns: 'Source Path' and 'Target Path'. The table has one row with 'C:\' in the 'Source Path' column and 'C:\indy\Sales data for Indy\C\' in the 'Target Path' column. At the bottom, there are two checkboxes: 'Start Mirror on connection' (checked) and 'Start Replication on connection' (checked). At the very bottom are three buttons: 'Connect', 'Cancel', and 'Help'.

Source Path	Target Path
C:\	C:\indy\Sales data for Indy\C\'

3. Verify that the **Start Replication on Connection** check box is selected. If replication is not initiated when the connection is established, data integrity cannot be guaranteed.

NOTE: Other tabs are available in the Connection Manager to set advanced connection settings. To establish a connection, you do not need to make modifications to these tabs; a connection can be established with the default settings. For more information on these advanced settings, see the chapter that corresponds with the tab name. For example, for information on the Verify tab, see [Verification](#) on page 9-1.

4. Click **Connect** when you are finished.

Controlling replication manually

To stop and start replication for an established source/target connection, right-click the connection on the right pane of the Management Console and select **Replication** and the appropriate replication control.

NOTE: If you stop replication manually, you will need to perform a remirror when it is restarted.

If you only need to temporarily pause data from being transmitted to the target and do not want to stop replication, select **Transmit, Pause**. When you are ready to restart the transmission, select **Transmit, Start**.


Monitoring replication

1. To view the connection on the Management Console, highlight the source machine on the left pane and the connection information will appear on the right pane. Use the horizontal scroll bar at the bottom of the right pane to view the **Replication Status** column. The possible states are:
 - ◆ **Replicating**—Data is being replicated to the target machine or replication data is in the queue waiting to be transmitted.
 - ◆ **Ready**—There is no data in the queue to replicate to the target machine.
 - ◆ **Stopped**—Replication has stopped.
 - ◆ **Pending**—If auto-remirror is enabled and you have experienced an auto-disconnect, the status will change to pending while the connections are reestablished and will update when the remirror begins. If auto-remirror is disabled and you have experienced an auto-disconnect, replication operations will queue on the source until a remirror is performed. At this point, the replication status will be pending. Without a remirror, data integrity cannot be guaranteed.
 - ◆ **Out of Memory**—The kernel has reached the stop replication limit or the machine has run out of physical memory. When out of memory occurs, all current active connections on the machine will show that a connection error has occurred. When the problem is fixed, the connections will switch to pending status.
2. Use the horizontal scroll bars to display additional replication statistics.
 - ◆ **Sent (Bytes)**—The total number of mirror and replication bytes that have been sent during this connection.
 - ◆ **Queued Replication (Bytes)**—The total number of replication bytes that remain in the source queue.
 - ◆ **Sent Replication (Bytes)**—The total number of replication bytes that have been sent during this connection.

NOTE: There are other statistics available through the Management Console, but they are not displayed by default. To change the statistics display, see [Customizing the statistics](#) on page 13-4.

Replicating permissions by name

1. To determine if Windows security is being replicated by name, right-click a source machine and select **Properties**.
2. Select the **Source** tab.



- ◆ If the **Replicate NT Security by Name** is marked, this feature is enabled and the local security ID is being replicated to the target. You will want to enable this feature if you are using local security groups.
 - ◆ If the **Replicate NT Security by Name** is not marked, this feature is disabled and domain security ID is being replicated to the target. You will want to disable this feature if you are using domain security groups.
3. Modify **Replicate NT Security by Name** if necessary.
 4. Click **OK** to save these settings.

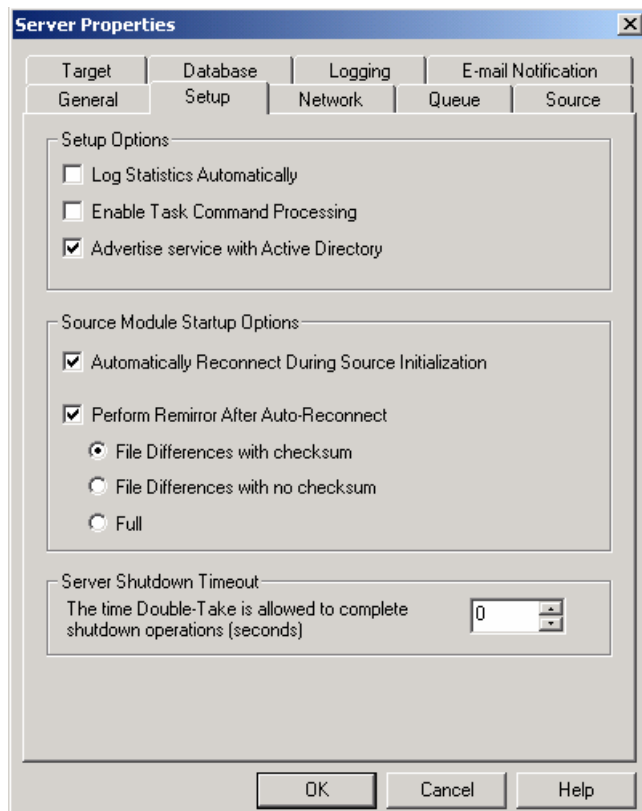
NOTE: Permissions cannot be set on a connection basis. If you modify this setting for a source, all connections made from that source machine will have this modified setting.

If you enable **Replicate NT Security by Name**, make sure that the same groups and users exist on the target as they do on the source. Otherwise, mirror and replication speed will be slower while adjustments are made for the missing groups or users.

Processing task commands

Task command processing is a Storage Mirroring feature that allows you to insert and run tasks at various points during the replication of data. Because the tasks are user-defined, you can achieve a wide variety of goals with this feature. For example, you might insert a task to create a snapshot or backup on the target after a certain segment of data from the source has been applied on the target.

1. To process a task, you must first enable task command processing. Right-click a source machine and select **Properties**.
2. Select the **Setup** tab.



3. By default, **Enable Task Command Processing** is disabled. Mark the field to enable it.
4. Since task command processing inserts a task inline with current replication, you must have an existing connection. Establish a connection according to [Establishing a new connection](#) on page 6-8, if you do not already have one established.
5. To insert tasks, use the DTCL `queuetask` command by following the text client instructions under [Processing task commands](#) on page 8-13. You may skip steps 1-3 in that section since you have already completed them through the Management Console.

Managing Replication Through the Text Client

For each of the replication commands, first identify which source you are working with by using the `source` command.

Command	<code>SOURCE</code>
Description	Identifies a machine as the active source machine
Syntax	<code>SOURCE <source_machine></code>
Options	<code>source_machine</code> —Name of the machine
Examples	<code>source indy</code>

Initiating replication on connection

1. If you do not know which replication set you will be connecting to the target machine, the `repset list` command will display the available replication sets for that source.

Command	<code>REPSET LIST</code>
Description	Lists all replication set names for the currently selected source
Syntax	<code>REPSET LIST</code>

2. Connect the replication set to the target by using the `connect` command. To initiate replication when the connection is established, specify the `replicate` switch with this command. By default, replication is started automatically.

Command	<code>CONNECT</code>
Description	Establishes a connection between a replication set and a target machine
Syntax	<code>CONNECT <repset> TO <target_machine></code> <code>MAP EXACT</code> <code> MAP BASE <target_path></code> <code> MAP <source_path> TO <target_path> [,...]</code> <code>[MIRROR NOMIRROR] [, REPLICATE NOREPLICATE] [, MONITOR NOMONITOR]</code> <code>[, ORPHANS NOORPHANS] [, COMPRESSION <level>] [CLEARRESTOREREQUIRED]</code>

Options

- ◆ **repset**—Name of the replication set
- ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address
- ◆ **MAP EXACT**—Specifies that the replication set data will be sent to the same logical volume on the target (c:\data and d:\files is copied to c:\data and d:\files, respectively)
- ◆ **MAP BASE target_path**—Substitute a complete path, including the volume, for **target_path** and the data will be replicated to target_path\SrcVolName on the target machine
- ◆ **MAP source_path TO target_path**—Custom location that specifies each directory on the source and where that data will be copied to on the target machine
- ◆ **...**—Indicates that the **source_path TO target_path** option can be used more than once for each source directory in the replication set
- ◆ **MIRROR**—Automatically initiates a mirror when the connection is established
- ◆ **NOMIRROR**—Does not initiate a mirror when the connection is established
- ◆ **REPLICATE**—Automatically initiates replication when the connection is established
- ◆ **NOREPLICATE**—Does not initiate replication when the connection is established
- ◆ **MONITOR**—Specifies that the target is going to monitor the specified source machine for failover. The source machine must have already been defined as a monitor machine.
- ◆ **NOMONITOR**—Specifies that the target is not going to monitor the source machine for failover
- ◆ **ORPHANS**—Removes orphan files on the target
- ◆ **NOORPHANS**—Does not remove orphan files on the target
- ◆ **COMPRESSION level**—Sets the compression level for data being sent to the target using levels 1 through 3 with 1 being minimum compression and 3 being maximum compression. This option sets the level only. It does not enable compression.
- ◆ **CLEARRESTOREREQUIRED**—Clears the restore required flag and initiates the connection

Examples

- ◆ connect Exchange to jersey map exact
- ◆ connect sql to jersey map base d:\DTFiles\

Notes

- ◆ The default settings for this command are mirror, replicate, nomonitor, and noorphans.
- ◆ If a path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.
- ◆ If you are establishing a connection within a NAT environment, you will need to specify the port of the router after the IP address (separated by a colon). See [NAT Environments](#) on page D-1 for more details.

Controlling replication manually

Identify the connection ID by locating the number in the left column below the target machine name and use one of the following commands for the appropriate replication control:

◆ Replication Start

Command	<code>REPLICATION START</code>
Description	Initiates the replication process
Syntax	<code><u>REPLICATION</u> START <conid *> [CLEARRESTOREREQUIRED]</code>
Options	<ul style="list-style-type: none">◆ conid—Connection ID assigned to the source/target connection◆ *—Specifies all connection IDs◆ CLEARRESTOREREQUIRED—Clears the restore required flag and initiates replication
Examples	<ul style="list-style-type: none">◆ <code>replication start 1</code>◆ <code>replication start *</code>

◆ Replication Stop

Command	<code>REPLICATION STOP</code>
Description	Stops the replication process
Syntax	<code><u>REPLICATION</u> STOP <conid *></code>
Options	<ul style="list-style-type: none">◆ conid—Connection ID assigned to the source/target connection◆ *—Specifies all connection IDs
Examples	<ul style="list-style-type: none">◆ <code>replication stop 1</code>◆ <code>replication stop *</code>
Notes	<p>If you stop replication manually, you will need to perform a remirror when it is restarted.</p> <p>If you only need to temporarily pause data from being transmitted to the target and do not want to stop replication, use the <code>transmission pause</code> command. When you are ready to restart the transmission, use the <code>transmission resume</code> command.</p>

Monitoring replication

Storage Mirroring offers a variety of methods for monitoring Storage Mirroring replication. Statistics include information like connection up-time, target status, bytes queued or sent, and so on. The informational data includes processing notifications, warnings or errors. Below is a brief description of how four of the monitoring methods are applicable to the Storage Mirroring connection. For detailed information on these and other monitoring methods, see [Monitoring Tools](#) on page 13-1.

- ◆ **Text Client**—The Storage Mirroring Text Client window displays a minimal number of statistical and informational data on screen.
- ◆ The DTCL `status` command displays statistical and informational data. Identify the connection ID by locating the number in the left column of the Text Client window and use the `status` command to display the information.

Command	<code>STATUS</code>
Description	Requests connection and statistical information
Syntax	<code>STATUS <u>CONNECT</u> MIRROR <u>REPLICATE</u> <u>TRANSMIT</u> <con_id></code>
Options	<ul style="list-style-type: none">◆ <code>CONNECT</code>—Displays connection information◆ <code>MIRROR</code>—Displays mirroring information◆ <code>REPLICATE</code>—Displays replication information◆ <code>TRANSMIT</code>—Displays transmission information◆ <code>con_id</code>—Connection ID assigned to the source/target connection
Examples	<ul style="list-style-type: none">◆ <code>status connect 1</code>◆ <code>status replicate 1</code>

- ◆ **LogViewer**—LogViewer is a utility within Storage Mirroring that gathers Storage Mirroring alerts and information messages. Alerts are processing notifications, warnings, and error messages. For detailed steps on how to use LogViewer, see [LogViewer](#) on page 13-39.
- ◆ **DTStat**—DTStats is a utility that takes snapshots of Storage Mirroring statistical data and writes the information to a file. For detailed steps on how to use DTStat, see [DTStat](#) on page 13-30.

Replicating permissions by name

1. To determine the permissions setting, use the `get` command with the `ReplicateNTSecurityByName` option.

Command	GET
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	GET <option> [<i>machine</i>]
Options	<ul style="list-style-type: none">◆ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.◆ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">◆ <code>get netport</code>◆ <code>get VerifyLogName</code>◆ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. To modify the permissions setting, use the `set` command with the `ReplicateNT SecurityByName` option. By default, this option is set to 0.
 - ◆ If `ReplicateNTSecurityByName` is enabled or set to 1, this feature is enabled and the local security ID is being replicated to the target. You will want to enable this feature if you are using local security groups.
 - ◆ If `ReplicateNTSecurityByName` is disabled or set to 0, this feature is disabled and the domain security ID is being replicated to the target. You will want to disable this feature if you are using domain security groups.

NOTE: Permissions cannot be set on a connection basis. If you modify this setting for a source, all connections made from that source machine will have this modified setting.

If you enable `ReplicateNTSecurityByName`, make sure that the same groups and users exist on the target as they do on the source. Otherwise, mirror and replication speed will be slower while adjustments are made for the missing groups or users.

Processing task commands

Task command processing is a Storage Mirroring feature that allows you to insert and run tasks at various points during the replication of data. Because the tasks are user-defined, you can achieve a wide variety of goals with this feature. For example, you might insert a task to create a snapshot or backup on the target after a certain segment of data from the source has been applied on the target.

To process a task, enable task command processing and then use the `queuetask` command.

1. Use the `get` command to view the current value of the `EnableTaskCmdProcessing` option. By default, the option is disabled and thus set to 0.

Command	GET
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	GET <option> [<i>machine</i>]
Options	◆ <i>option</i> —See Storage Mirroring Text Client program settings on page 14-15. ◆ <i>machine</i> —Name of the machine
Examples	◆ <code>get netport</code> ◆ <code>get VerifyLogName</code> ◆ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. Use the `set` command to enable task command processing by setting `EnableTaskCmdProcessing` to 1.

Command	SET
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	SET <option>=<value> [<i>machine</i>]
Options	◆ <i>option</i> —See Storage Mirroring Text Client program settings on page 14-15. ◆ <i>value</i> —See Storage Mirroring Text Client program settings on page 14-15. ◆ <i>machine</i> —Name of the machine
Examples	◆ <code>set netport=1100</code> ◆ <code>set VerifyLogName="HA server.log"</code> ◆ <code>set DefaultAddress="129.12.34.3"</code>
Notes	◆ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed. ◆ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

3. Since task command processing inserts a task inline with current replication, you must have an existing connection. Establish a connection according to [Establishing a new connection](#) on page 6-16, if you do not already have one established.

-
4. Use the `queuetask` command to places task(s) in the replication queue inline with current processing.

Command `QUEUETASK`

Description Queues tasks inline with replication data.

Syntax `QUEUETASK <job_name> TO <target_machine> ONQUEUE = <task> [args]
 | ONTRANSMIT = <task> [args] | ONRECEIVE = <task> [args] | ONEXECUTE = <task> [args]
 [TIMEOUT = <timeout>] [INTERACT | NOINTERACT]`

Options

- ◆ **job_name**—Any name that you have assigned to the job. This will be the identifier you will see in the log files.
- ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name. The target is required even if you are only queuing a task to be executed on the source.
- ◆ **ONQUEUE**—Execute the specified task on the source machine as soon as the source receives and queues the task. During heavy replication, there may be a delay while the task is queued inline with the replication operations.
- ◆ **ONTRANSMIT**—Execute the specified task on the source machine just before the source transmits the task to the target.
- ◆ **ONRECEIVE**—Execute the specified task on the target machine as soon as the target receives and queues the task.
- ◆ **ONEXECUTE**—Execute the specified task on the target when the target processes the task from the queue. Since the task is not executed until it is processed, if the target is paused, the task will be held in queue.
- ◆ **task**—The path and filename of the task to run relative to the machine it will be run on. Tasks include any valid executable or batch file. If the path contains non-alphanumeric characters, you must enclose it in quotation marks. The executables or files must exist in the specified location on the machine where they will be executed
- ◆ **args**—Arguments or options which need to be supplied with the task. Multiple arguments can be supplied in a space-separated list enclosed in quotation marks.
- ◆ **TIMEOUT**—Length of time to pause while waiting for the task to complete
- ◆ **timeout**—Any valid number followed by an optional time indicator indicating length of time. The valid time indicators include seconds, minutes, hours, and days. If you do not specify a time indicator, it will default to seconds. The number zero (0) indicates there is no timeout delay and the next operation is immediately processed. The keyword `FOREVER` indicates that the next operation is not processed until the task has completed execution. No `TIMEOUT <timeout>` setting will default to forever.
- ◆ **INTERACT**—Tasks interact with the desktop and, therefore, display on screen and run in the foreground
- ◆ **NOINTERACT**—Tasks do not interact with the desktop

Examples `qtask backup to jersey onrec=PauseAndBackup.bat onexec=resume.bat`

Notes

- ◆ The default setting for this command is `nointeract`.
- ◆ Any combination of one or more execution points can be used with the same `queuetask` command.
- ◆ All processing messages, including errors, can be viewed in the Storage Mirroring logs and the Windows Event Viewer.
- ◆ `Onqueue` will still execute as soon as the task is placed on the queue even if transmission is stopped (manually stopped or paused, unmet scheduled transmission criteria, etc.). Any other option will not execute until transmission is restarted.
- ◆ If your source is in a restore required state, any task placed on the queue will be executed immediately. Use caution when submitting tasks while in this state so that the target does not get inadvertently updated.
- ◆ A task may be discarded in the following scenarios: all connections to a target are manually disconnected, replication is stopped for all connections to a target, or an auto-disconnect occurs.
- ◆ If a task is submitted after replication is stopped, the task will be executed immediately.
- ◆ If you disable task command processing while tasks are in queue, those tasks will not be executed.
- ◆ The user submitting the task command must be a member of the **Double-Take Admin** security group on both the source and target and the Storage Mirroring service must have proper privileges to access the files or run the commands specified in the task.

9 ► Verification

Verification is the process of confirming that the data on the target is identical to the data on the source. Verification creates a log file detailing what was verified as well as which files are not synchronized. If the data is not the same, Storage Mirroring can automatically initiate a remirror. The remirror ensures data integrity between the source and target.

NOTE: Because of the way the Windows Cache Manager handles memory, machines that are doing minimal or light processing may have file operations that remain in the cache until additional operations flush them out. This may make Storage Mirroring files on the target appear as if they are not synchronized. When the Windows Cache Manager releases the operations in the cache on the source, the files will be updated on the target.

Verification Options

In the Storage Mirroring Management Console, verification can be configured when a connection is established or initiated manually at any time. (Verification cannot be initiated while a mirror is occurring.) Verification can also be scheduled to occur at periodic intervals. In the Storage Mirroring Text Client, verification can only be initiated after a connection is established. In either case, you must specify how you want the Storage Mirroring verification process to execute. The following options can be set:

- ◆ **Start Time**—The start time specifies the time for the verification process to begin.
- ◆ **Verification Interval**—The verification interval specifies how often the verification process will be repeated. This interval can be as short as minutes or as long as days.
- ◆ **Remirror**—You can configure Storage Mirroring to automatically perform a remirror for only those files that are not synchronized on the source and target. If the verify process indicates that the files are identical on both machines, the remirror is not performed. There are two remirror options available when verification has found inconsistencies:
 - ◆ **Only if newer on source**—If a file on the source is newer than the file on the target based on date and time stamps, the file is mirrored to the target machine. This option overwrites the file on the target machine. If the file on the target is newer, the file on the source is not mirrored to the target.

NOTE:	Accessing files on the target, directly or indirectly, could cause them to appear newer on the target. The Only if Newer on Source option will prevent newer files on the target from being successfully overwritten. For example, when a database application is started on the target, the date and time stamp on the file will be updated, although no data may actually change. If the Only if Newer on Source option is used, the database file on the target in this example, will not be updated, even though the source file may have newer data.
--------------	---

- ◆ **Block checksum**—With this option, each file that fails the verification process is read in blocks to determine which blocks are not identical. Those blocks that are different are sent to the target machine. With this option, the entire file is not overwritten on the target machine; just the dissimilar blocks.

The following table illustrates how the remirror options work together, as well as how they work with the global checksum setting on the Source tab of the Server Properties. Determine which options you need to select when configuring verification by identifying which type of action you want Storage Mirroring to perform during the verification.

NOTE:	Verification does not identify files that exist on the target that are no longer on the source machine. See orphan files for information on deleting files on the target that no longer exist in the source replication set.
--------------	--

Remirror Options When Verifying

Block Checksum All Server Properties Source Tab ^a	Remirror Automatical ly Verify Option	Only if Source is Newer Verify Option	Block Checksum Verify Option	Action Performed
(✓)	✓			Any file that is different on the source and target based on the date, time, and/or size is transmitted to the target. The mirror sends the entire file.
(✓)	✓	✓		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	✓		✓	Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
✓	✓		✓	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(✓)	✓	✓	✓	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

a. A checkmark enclosed in parentheses (✓) indicates that the global option on the Server Properties source tab can be on or off. The use of this option does not change the action performed during the verification.

NOTE: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the Block Checksum All option on the Server Properties Source tab to ensure proper file comparisons.

Verification Log

A verification log is created on the source during the verification process. The log identifies what is verified as well as which files are not synchronized. The options and defaults are listed below.

- ◆ **File Name and Location**—By default, the verification log is called `DTVerify.log` and is stored in the same directory as the Storage Mirroring program files.
- ◆ **Append or Overwrite**—By default, the log is appended each time a verification process is completed. The log file clearly displays which source, replication set, and connection the verification was performed on so that multiple verification processes can be easily identified.
- ◆ **File Size Limit**—By default, the log is limited to a file size of 1 MB. When the log file reaches this limit, no additional data will be logged.

Sample Verification Log

```
----- VERIFICATION OF CONNECTION 2 (Sales data for Indy --> 206.31.65.40 : 1100) -----
Start Time: 5/24/2005 12:15:20 PM for connection 2 (Sales data for Indy --> 206.31.65.40 : 1100)
File:      jersey\users\bob\budget.xls  DIFFERENT ON TARGET
          Source Attributes:  Timestamp = 1/17/2005 8:21:36 PM  Size = 1272  Mask = [0x20]
          Target Attributes:  Timestamp = 1/17/2005 8:21:36 PM  Size = 1272  Mask = [0x20]
Security descriptors are different.
0 BYTES OUT OF SYNC

File:      jersey\users\bill\timesheet.xls  DIFFERENT ON TARGET
          Source Attributes:  Timestamp = 1/17/2005 8:21:37 PM  Size = 1272  Mask = [0x20]
          Target Attributes:  Timestamp = 1/17/2005 8:21:37 PM  Size = 1272  Mask = [0x23]
0 BYTES OUT OF SYNC

File:      jersey\users\vincent\training.doc  DIFFERENT ON TARGET
          Source Attributes:  Timestamp = 3/12/2005 3:28:20 PM  Size = 17  Mask = [0x20]
          Target Attributes:  Timestamp = 5/24/2005 5:05:26 PM  Size = 2  Mask = [0x20]
17 BYTES OUT OF SYNC

Completion Time: 5/24/2005 12:15:20 PM for connection 2 (Sales data for Indy --> 206.31.65.40 : 1100)
Elapsed Time (seconds): 0.256470
Total Directories Compared: 6
Total Directories Missing: 0
Total Directories Remirrored: 0
Total Files Compared: 120
Total Files Missing: 0
Total Files Different: 3
Total Files Encrypted: 0
Total Files Remirrored: 1
Total Bytes Skipped: 18812506
Total Bytes Compared: 185272
Total Bytes Missing: 0
Total Bytes Different: 0
Total Bytes Remirrored: 0
Related links and directory attributes have been adjusted.
----- END OF VERIFICATION -----
```

The verification process is delineated by beginning and end markers. These separate different verification process. A list of files that are different on the source and target is provided as well cumulative totals for the verification process. The information provided for each file is the state of its synchronization between the source and the target at the time the file is verified. If the remirror option is selected so that files that are different are remirrored, the data in the verify log reflects the state of the file before it is remirrored, and does not report the state of the file after it is remirrored. If a file is reported as different, review the output for the file to determine what is different.

- ◆ **Timestamp**—The last modified date and time of the file
- ◆ **Size**—The size, in bytes, of the file
- ◆ **Mask**—The attributes associated with the file (additional details on the following page)
- ◆ **Security descriptors**—The NTFS file permissions of the file. If the file permissions are different, the message "Security descriptors are different" will be logged. If the file permissions are the same, nothing will be logged.

- ◆ **Bytes out of sync**—The number of bytes that are not synchronized between the file on the source and the file on the target. If the data in the file is identical, the message "0 BYTES OUT OF SYNC" will be logged. If the file is different, the message will indicate how many bytes were different. This message does not indicate that the file was remirrored during the verify.

The mask must be converted in order to determine what attributes are assigned to a file. The mask is a hexadecimal number corresponding to a binary number that indicates what the attributes are. Using the following steps, you can determine how the mask corresponds to the attributes of a file.

1. Each mask begins with 0x. Identify the hexadecimal number after the constant 0x. For example, if the mask is 0x23, then the hexadecimal number you are interested in is 23. The hexadecimal number may be up to four digits.
2. Convert the hexadecimal number to its 16-digit binary equivalent. You can use the Windows calculator for this conversion.
 - a. Select **Start, Programs, Accessories, Calculator**.
 - b. Switch to scientific view, if it is not already in that view, by selecting **View, Scientific**.
 - c. Select **Hex**.
 - d. Enter the hexadecimal number, for example 23, as specified in your verification log.
 - e. Select **Bin** and the hexadecimal number will change to the binary equivalent.
 - f. Pad the beginning of the binary equivalent with zeroes (0) so that the number is 16 digits long. For example, hexadecimal number 23 converts to 100011, so the 16-digit binary equivalent would be 0000000000100011.
3. Determine what number appears in each position of the binary number. Because binary numbers count from right to left, start with position 1 on the right. For example, hexadecimal number 23, which converted to 0000000000100011, is broken out into the following positions.


Position (from right to left)	Sample Hexadecimal Number 23	Enter Your Hexadecimal Number	Attribute
1	1		Read only
2	1		Hidden
3	0		None
4	0		System
5	0		Directory
6	1		Archive
7	0		Encrypted
8	0		Normal
9	0		Temporary
10	0		Sparse file
11	0		Reparse point
12	0		Compressed
13	0		Offline
14	0		Not content indexed
15	0		None
16	0		None

4. Using the same chart above, identify those attributes that are enabled by those positions equal to one (1). The positions equal to zero (0) are disabled and that attribute does not apply. So hexadecimal number 23, which converted to 0000000000100011, indicates read only, hidden, and archive. Another example might be mask 0x827 which converted to binary is 0000100000100111. Positions 1-3, 6, and 12 are all enabled which indicates the file is read only, hidden, archive, and compressed.

NOTE: Files that were replicated with the Replicate NT Security by Name feature enabled, will be identified as different in the log file because of the local name attribute. The files will be the same.

How Verification Works

To better understand how Storage Mirroring verification works, the following diagram identifies the steps involved with the Storage Mirroring verification process. Each numbered step is described below the diagram.



1. Verification is initiated by the Storage Mirroring administrator, either manually through one of the clients or automatically through a scheduled verify.
2. Storage Mirroring compares the date, time, and size of files on the source to the files on the target and determines which files are not identical. Also, you can select the option to have Storage Mirroring perform a checksum calculation on all Storage Mirroring files regardless of the date, time, and size.
3. A verification log file is created which details what was verified as well as which files are not synchronized.
4. If a remirror is specified, Storage Mirroring transmits only the dissimilar data by the remirror options chosen, newer on source or checksum.
5. As each packet of mirror data is received on the target, the target returns an acknowledgment to the source confirming that the mirrored data has been received.

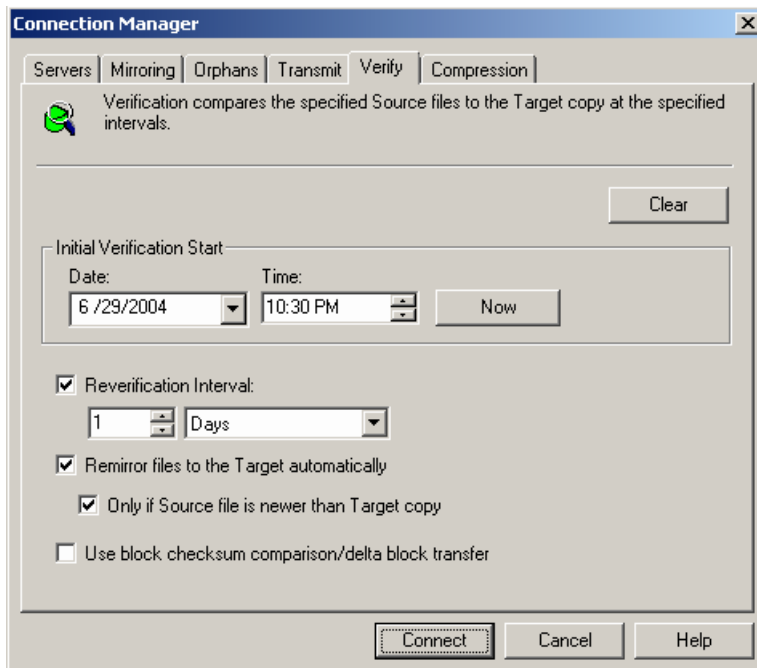
Managing Verification With the Management Console

In the Management Console, verification can be configured when a connection is established or it can be initiated manually at any time while a connection is active.

Configuring verification on connection

1. There are four methods available for establishing a connection and opening the Connection Manager:
 - ◆ Highlight a replication set and select **Tools, Connection Manager**.
 - ◆ Right-click on a replication set and select **Connection Manager**.
 - ◆ Drag and drop a replication set to a target machine on the left pane of the Management Console.
 - ◆ Highlight a replication set and then drag and drop the replication set to a target machine on the right pane of the Management Console.

The Connection Manager will open and the Servers tab will be selected. Complete the Servers tab entries as described in [Establishing a new connection](#) on page 6-8 .
2. Select the **Verify** tab.

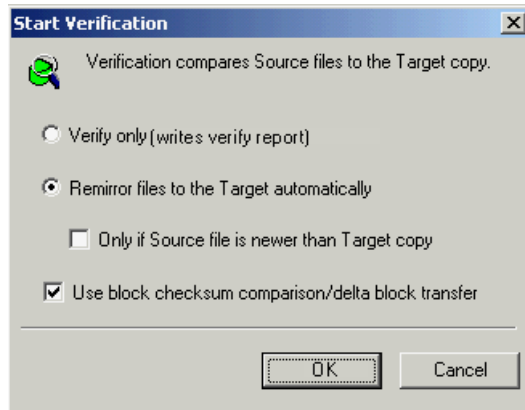


3. Select the verification options that you would like to perform.
 - ◆ The **Initial Verification Start** field allows you to select the immediate date and time by clicking **Now** or entering in a specific **Date** and **Time**. The down arrow next to **Date** displays a calendar allowing easy selection of any date. **Time** is formatted for any AM or PM time.
 - ◆ Mark the **Reverification Interval** check box to repeat the verification process at the specified interval. Specify an amount of time and choose minutes, hours, or days.
 - ◆ Mark the **Remirror files to the Target automatically**, **Only if Source file is newer than Target copy**, and/or **Use block checksum comparison/delta block transfer** check boxes to perform the type of remirror you desire. To determine the options to select for the type of remirror you wish to perform, see the chart [Remirror Options When Verifying](#) on page 9-3.

NOTE: The settings on the Verify tab are dependent on some of the settings on the Server tab. If any of the top four Server settings (**Source Server**, **Replication Set**, **Target Server**, **Route**) are modified after verification settings are established, the verification settings will be cleared and will need to be reset.

Initiating verification manually

1. To initiate verification manually, right-click the connection on the right pane of the Management Console and select **Verify**.



2. Select the verification options that you would like to perform.
 - ◆ **Verify Only**—This option performs the verification process by comparing the date, time and size of each file and generates a verification report identifying the files that are not synchronized.
 - ◆ **Remirror Options**—Mark the **Remirror files to the Target automatically**, **Only if Source file is newer than Target copy**, and/or **Use block checksum comparison/delta block transfer** check boxes to perform the type of remirror you desire. To determine the options to select for the type of remirror you wish to perform, see the chart [Remirror Options When Verifying](#) on page 9-3 .

Verification log

A verification log is created during the verification process. The log details what is verified as well as which files are not synchronized. Use the following steps to set up your verification log in the Management Console.

NOTE: All changes to a verification log should be made on the source side of a connection.

1. Right-click a source machine on the left pane of the Management Console and select **Properties**.
2. Select the **Logging** tab.

The screenshot shows the 'Server Properties' dialog box with the 'Logging' tab selected. The 'Folder' field is set to 'C:\Program Files\OpenView\Storage Mirroring'. Under the 'Messages & Alerts' section, 'Maximum Length (bytes)' is 1048576 and 'Maximum Files' is 5. Under the 'Verification' section, 'Filename' is 'DTVerify.log', 'Maximum Length (bytes)' is 1048576, the 'Append' checkbox is checked, and 'Language' is 'English'. Under the 'Statistics' section, 'Filename' is 'statistic.sts', 'Maximum Length (bytes)' is 10485760, and 'Write Interval (minutes)' is 5. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

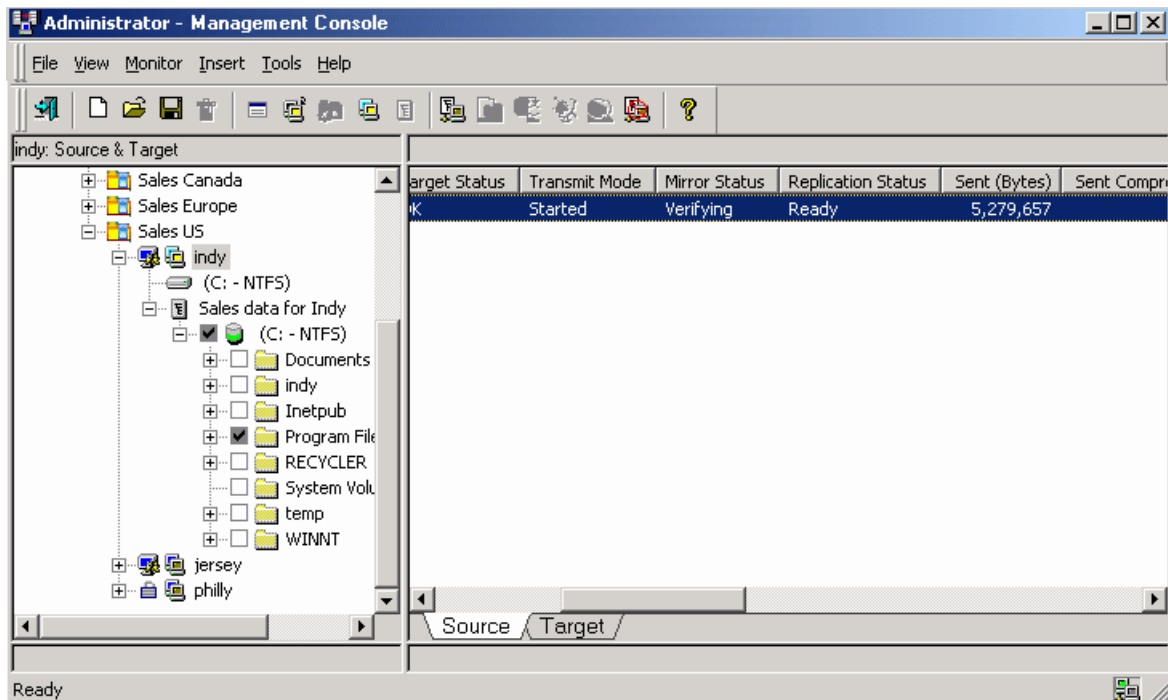
3. At the top of the window, **Folder** identifies the location where the log files identified on this tab are stored.
4. Under the Verification section, **Filename** contains the log file for the verification process. Update this field, if necessary. The default is `DTVerify.log`.
5. Specify the **Maximum Length** of the log file. The default is 1048576 bytes (1MB).
6. By default, the log is appended each time a verification process is completed. Clear the **Append** check box if you do not want to append to the previous log file.

NOTE: Changes made to the verification log in the **Server Properties, Logging** tab will apply to all connections from the current source machine.

Monitoring verification

To view the connection on the Management Console, highlight the source machine on the left pane and the connection will appear on the right pane. Use the horizontal scroll bar at the bottom of the right pane to view the **Mirror Status** column; the status of the mirroring operation for this connection is displayed. The possible states are

- ◆ **Idle**—No mirroring or verifying activity is taking place.
- ◆ **Mirroring**—If the size of the replication set was not calculated prior to the connection being established, the status indicates **Mirroring**.
- ◆ **Percentage Complete**—If the size of the replication set was calculated prior to the connection being established, the percentage of the mirror that is complete will be displayed.
- ◆ **Paused**—No new mirror data is being written to the mirror queue. The target is busy and sent a request to the source to pause while the target frees its memory.
- ◆ **Waiting**—The source has finished sending the data to the target, but the target has not finished writing to the disk.
- ◆ **Verifying**—The source is verifying if the data on the source and target are identical.
- ◆ **Removing Orphans**—Storage Mirroring is checking for orphan files within the target path location that exist on the target but not on the source. These files will be removed.



Managing Verification Using the Text Client

In the Storage Mirroring Text Client, verification can only be initiated after a connection is established.

Initiating verification

1. Log on to the source using the `login` command.

Command	LOGIN
Description	Log on to a Storage Mirroring machine
Syntax	LOGIN <machine> <username> <password> [domain]
Options	<ul style="list-style-type: none">♦ machine—Name of the machine♦ username—Name of the user♦ password—Password associated with username.♦ domain—If logging in using a domain account, this is the domain name. If logging in using a local account, this is the machine name.
Examples	login indy administrator *****
Notes	<ul style="list-style-type: none">♦ The <code>login</code> command is not available when scrolling through the Text Client command history.♦ If characters in the password include non-alphanumeric characters, the password field must be enclosed in quotation marks.♦ The password cannot be a Storage Mirroring keyword. These are any DTCL command (source, target, and so on.) or any DTCL shortcut command (env, mon, rep, and so on).

2. Identify the source that you want to initiate verification for by using the `source` command.

Command	SOURCE
Description	Identifies a machine as the active source machine
Syntax	SOURCE <source_machine>
Options	source_machine —Name of the machine
Examples	source indy

3. Identify the connection ID by locating the number in the left column below the target machine name and initiate the verification process by using the `verify` command.

Command	VERIFY
Description	Verifies the integrity of the data between the source and target machines.
Syntax	VERIFY <con_id> [SYNC [,NEWER][,CHECKSUM NOCHECKSUM] NOSYNC] [,CHECKSUM NOCHECKSUM] [STARTTIME=<mm/dd/yy> [hh:mm]] [EVERY <number> <time_units>] [ORPHANS NOORPHANS] [CLEARRESTOREREQUIRED]

Options

- ◆ **con_id**—The connection ID which should be verified
- ◆ **SYNC**—Synchronizes any data that is different
- ◆ **NEWER**—Mirrors only those files that are newer on the source than on the target
- ◆ **CHECKSUM**—Mirrors those blocks that are different on the source than on the target based on checksum comparisons
- ◆ **NOCHECKSUM**—Does not perform a checksum comparison when mirroring files
- ◆ **NOSYNC**—Do not synchronize any data that is different
- ◆ **STARTTIME**—Starts the verification process at the time specified
- ◆ **mm/dd/yy**—Date in month/day/year format when the verification process will begin
- ◆ **hh:mm**—Time in hour:minute format using the 24 hour clock when the verification process will begin
- ◆ **EVERY**—Repeat the verification process at the frequency specified
- ◆ **number**—Length of time to repeat the verification process
- ◆ **time_units**—Minutes (**min**), hours (**hr**), or days (**day**)
- ◆ **ORPHANS**—Removes orphan files on the target
- ◆ **NOORPHANS**—Does not remove orphan files on the target
- ◆ **CLEARRESTOREREQUIRED**—Clears the restore required flag and initiates the verification

Examples

- ◆ `verify 1`
- ◆ `verify 2 sync, newer`
- ◆ `verify 2 every 2 hr`

Notes

The default verification settings are sync, checksum, and noorphans.

Setting the verification log attributes

1. To determine the current status of the verification log, use the `get` command with the `VerifyLogName` option.

Command	<code>GET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	<code>GET <option> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">◆ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.◆ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">◆ <code>get netport</code>◆ <code>get VerifyLogName</code>◆ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

```
Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)

indy:
  verifylogname = DTVerify.log
  (press any key)

> connect marketing to jersey map exact mirror,replicate
> verify 1 sync ,newer every 5 hr
> get verifylogname
```

2. To modify the verify log file name, use the `set` command with the `VerifyLogName` option. The `VerifyLogName` option is any valid file name.

Command	<code>SET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	<code>SET <option>=<value> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">◆ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.◆ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.◆ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">◆ <code>set netport=1100</code>◆ <code>set VerifyLogName="HA server.log"</code>◆ <code>set DefaultAddress="129.12.34.3"</code>

Notes

- ◆ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to ClientLogName to take effect. And the service must be restarted to cause a change in the modules loaded if the LoadSourceTarget setting is changed.
- ◆ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

3. To determine the current status of the verification log append setting, use the `get` command with the `VerifyLogAppend` option.
4. To modify the verify log file append setting, use the `set` command with the `VerifyLogAppend` option. The `VerifyLogAppend` option is set to 1 for append or 0 for overwrite. By default, log file is appended.
5. To determine the maximum file size of the verification log, use the `get` command with the `VerifyLogLimit` option.
6. To modify the maximum file size of the verification log in bytes, use the `set` command with the `VerifyLogLimit` option. The default size is 1048576 bytes and is limited by the available hard drive space up to 4194304 KB (4 GB).

Monitoring verification

In addition to the verification log, Storage Mirroring offers a variety of methods for monitoring Storage Mirroring activity. For detailed information, see [Monitoring Tools](#) on page 13-1.

The Storage Mirroring replication and mirroring processes place replication set data on queues on the source and the transmission process sends the data from the source to the target. There are three methods for transmitting the Storage Mirroring data:

- ◆ **Real-time transmission**—Storage Mirroring data is continuously transmitted to the target machine. Although the data may be queued if the network or target machine is slow, the data is transmitted as soon as possible. This is the default transmission setting.
- ◆ **Scheduled transmission**—Storage Mirroring data is transmitted to the target machine when event driven or scheduling criteria is met. Storage Mirroring maintains the source/target connection as seen in the Storage Mirroring clients, but the transmission of data is delayed until the event driven or scheduling criteria is met. The mirror and replication data is queued until the transmission begins.
- ◆ **Manual transmission**—Storage Mirroring data is transmitted to the target when the network administrator manually initiates the transmission. Storage Mirroring maintains the source/target connection as seen in the Storage Mirroring clients, but the transmission of data is delayed until the network administrator manually initiates the process. The mirror and replication data is queued until the transmission begins.

For each type of transmission, real-time, scheduled, or manual, bandwidth limitations are available to restrict the amount of network bandwidth used for Storage Mirroring data transmissions.

NOTE: You can also pause the execution of operations on the target rather than modifying how the data is transmitted from the source.

Bandwidth Limitations

For each type of transmission, real-time, scheduled, or manual, bandwidth limitations are available to restrict the amount of network bandwidth used for Storage Mirroring data transmissions. The network administrator specifies a percentage of bandwidth that is available or an absolute bandwidth limit for Storage Mirroring transmissions and Storage Mirroring never exceeds that allotted amount. The bandwidth not in use by Storage Mirroring is available for all other network traffic.

Real-Time and Manual Transmission

Real-time transmission is the default Storage Mirroring setting. When Storage Mirroring is mirroring and replicating, data is continuously sent to the target machine. The only way to stop the transmission of data is to manually intervene. Manual intervention would again be required to restart the transmission.

NOTE: You can stop or pause the transmission of data to the target machine. In both cases, Storage Mirroring queues data on the source. The difference between these two settings is the effect on the network connection. When you stop the transmission of data, the network connection between the source and target is terminated and a new connection will have to be made to restart the transmission. When you pause the transmission of data, the network connection between the source and target is maintained and can be restarted with the same connection information as when the pause was made.

Scheduled Transmission

Mirror and replication data can be held in a queue until scheduling criteria that you establish has been met. When any of the scheduling options are met, transmission to the target is started. The transmission can also be stopped by using scheduling criteria.

Scheduled transmission options can be toggled on and off, allowing you to enable them only when you need to use them.

NOTE: Storage Mirroring checks the schedule once every second, and if a user-defined criteria is met, transmission will start or stop, depending on the option specified.

Scheduled Transmission Options

The following transmission options are available through the Management Console and the Text Clients. They are user configurable settings which can occur in any order or combination. The options can be set and then disabled, allowing you to enable them only when you need to use them.

NOTE: The transmission options are set for each target machine and cannot be defined separately for each replication set or source/target combination. This means that all connections to a target will have the same transmission options despite the source and replication set. For example, if you have two replication sets on the same source and these replication sets are connected to the same target, any scheduling criteria established for one connection will also be applied to the other connection.

- ◆ **Storage Mirroring Start Transmission Criteria**—When Storage Mirroring is not transmitting data, it places the information in a queue on the source. When transmission begins, Storage Mirroring sends all of the data from the queue to the target machine in the same order it was placed on the queue. Through the transmission start criteria you can set four triggers which initiate the transmission of data from the source to the target.
 - ◆ **Transmission Session Start**—The transmission session start establishes a date and time of the day to begin transmitting data. For example, you may want to specify a transmission time that corresponds to a low bandwidth usage time. Once started, Storage Mirroring will continue to transmit data until the queue is empty or until another limitation stops the transmission.
 - ◆ **Session Interval**—The session interval begins transmitting Storage Mirroring data at specified intervals of time. This option is used in conjunction with **Transmission Session Start**. For example, if the **Session Interval** is set to repeat transmission every 30 minutes and the **Transmission Session Start** is set to begin transmitting at 10 p.m., if the queue is emptied at 10:20 the transmission will stop. The start criteria is again met at 10:30 and Storage Mirroring will begin transmitting any new data in the queue.
 - ◆ **Queue Thresholds**—If the allocated amount of queue disk space is in use, Storage Mirroring cannot continue to queue data causing an auto-disconnect and the potential for loss of data. To avoid using the entire queue, you can configure Storage Mirroring to begin transmitting data to the target when the queue reaches a certain point. This point can be defined as a percentage of the disk queue that must be in use or the number of bytes in the disk queue. For example, if you specify 40%, when 40% of the queue is in use, Storage Mirroring initiates the transmission process and sends the data in the queue to the target machine. The transmission stops when the queue is empty or a Storage Mirroring stop transmission criteria is met. Or you might set a queue threshold of 500 MB. Storage Mirroring will wait until there is 500 MB of data in the queue and then begin transmitting the data. Like other start criteria, Storage Mirroring continues transmitting until the queue is empty or a Storage Mirroring stop criteria is met.

NOTE: A **Transmission Session Start** setting will override any other start criteria. For example, if you set the **Transmission Session Start** and the **Queue Threshold**, transmission will not start until you reach the indicated start time.

- ◆ **Storage Mirroring Stop Transmission Criteria**—There may be instances when Storage Mirroring is currently transmitting data, but the transmission should be stopped. This might occur at high traffic times when transmitting Storage Mirroring data might affect network traffic. Through the transmission stop criteria you can set two triggers which stop the transmission of data from the source to the target.

NOTE: Storage Mirroring checks the schedule once every second, and if a user defined criteria is met, transmission will stop.

- ◆ **Time Limit**—The time limit specifies the maximum length of time for each transmission period. Any data that is not sent during the specified time limit remains on the source queue. When used in conjunction with the session interval start option, you can explicitly define how often data is transmitted and how long each transmission lasts.
- ◆ **Byte Limit**—The byte limit specifies the maximum number of bytes that can be sent before ending the transmission session. When the byte limit is met, Storage Mirroring will automatically stop transmitting data to the target. Any data that still remains waits in the source queue until the transmission is restarted. When used in conjunction with a session start option, you can explicitly define how much data is being sent at a given time.

NOTE: The transmission start and stop criteria should be used in conjunction with each other. For example, if you set the Queue Threshold equal to 10 MB and the Byte Limit equal to 10 MB, a network connection will be established when there is 10 MB of data in the queue. The data will be transmitted and when the 10 MB Byte Limit is reached, the network connection closes. This is useful in configurations where metered charges are based on connection time.

-
- ◆ **Setting a Storage Mirroring Transmission Window**—The transmission window establishes a period of availability for all Storage Mirroring transmissions. You can specify a begin and end time for all Storage Mirroring transmissions. When a transmission window is in effect, all other start and stop criteria are bound by this window. This means that Storage Mirroring will never transmit data outside of an established window, regardless of other transmission settings.

For example, if you set a window of availability from 9 p.m. to 4 a.m. and a start option to initiate transmission at 5 a.m., the window option will override the start option and no data will be sent at 5 a.m.

NOTE: Setting a transmission window by itself is not sufficient to start a transmission. You still need to set a start criteria within the window.

Setting Transmission Criteria Using the Management Console

Transmission criteria allows the user to specify when a transmission will take place. These settings can either be set before or after a connection is made, and while mirroring and replication are being performed.

Establishing transmission criteria when configuring a connection

After you have configured a source/target connection by following the steps in [Establishing a new connection](#) on page 6-8, you can specify any transmit options by selecting the Transmit tab from the Connection Manager.

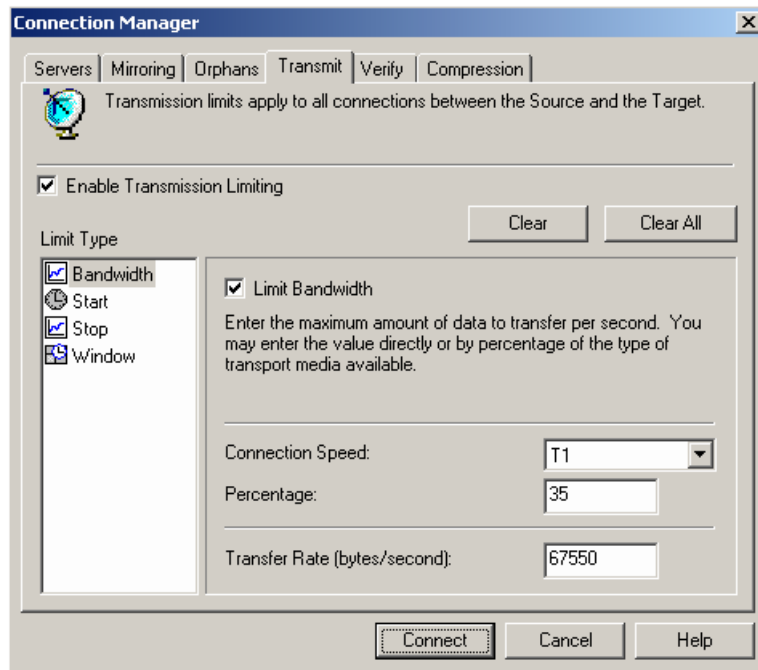
The **Transmit** tab contains four limit types: Bandwidth, Start, Stop, and Window. The transmission options for each limit type are displayed by highlighting a selection in the **Limit Type** box.

At the top of the **Transmit** tab dialog box, the **Enable Transmission Limiting** check box allows you to turn the transmission options on or off. You can enable the transmission options by marking the **Enable Transmission Limiting** check box when you want the options to be applied, but you can disable the transmission options, without losing the settings, by clearing that check box.

Also at the top of the **Transmit** tab dialog box, the **Clear All** button, when selected, will remove all transmission limitations that have been set under any of the four limit types. The **Clear** button will clear the settings for the **Limit Type** selected.

Setting bandwidth limitations

1. Select the **Bandwidth** option in the **Limit Type** box.



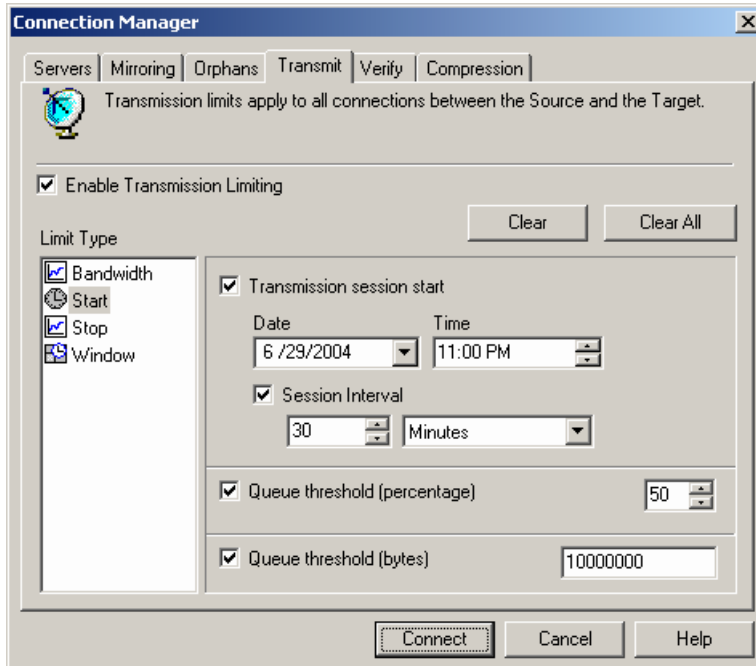
2. Mark the **Limit Bandwidth** check box to enable the bandwidth limiting features.
3. Define the bandwidth available for Storage Mirroring transmission
 - ◆ **Connection Speed**—Specify the total bandwidth capacity that is available
 - ◆ **Percentage**—Specify the percentage of bandwidth to be used for Storage Mirroring transmissionsThe **Transfer Rate** will automatically be calculated.
4. If no additional transmission options need to be set, establish the connection by clicking **Connect**.

NOTE: The only value that is persistently stored is the **Transfer Rate**. When the page is refreshed, the percentage and available bandwidth may not be the same value that you entered. Storage Mirroring changes these values to the maximum values for the smallest possible link. For example, if you enter 10% of T1, this is equivalent to 19300 bytes/second. If you select another transmit option or another tab of the Connection Manager and then return to the bandwidth option, Storage Mirroring will change the value to 60% of 256 Kbps which is also equivalent to 19300 bytes/second.

Setting transmission start criteria

Transmission start criteria initiates Storage Mirroring data transmission which continues until the queue is empty or a transmission stop criteria is met.

1. Select the **Start** option in the **Limit Type** box.

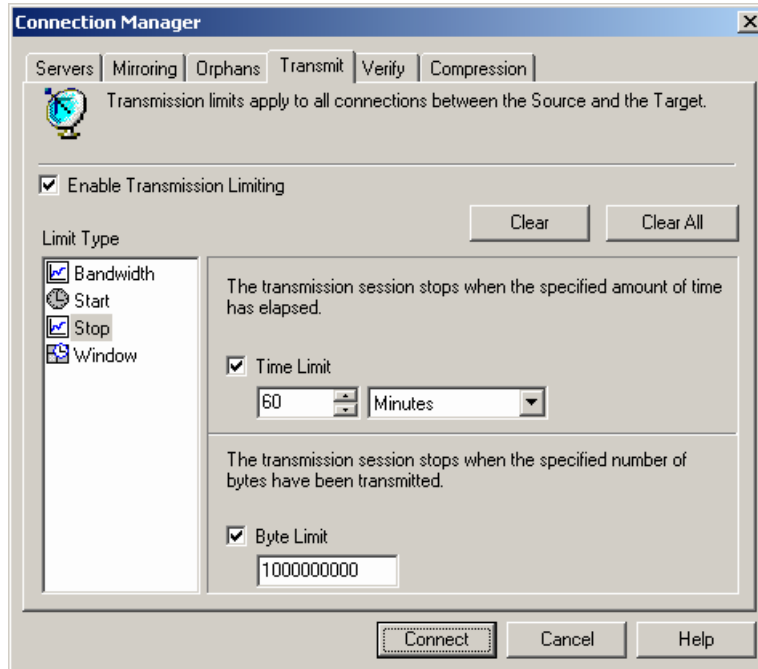


2. Define the start options for Storage Mirroring transmission by using any combination of the following options:
 - ◆ **Transmission session start**—Specify a **Date** and **Time** to start transmitting data. The down arrow next to the date field displays a calendar allowing easy selection of any date. The time field is formatted for any AM or PM time.
 - ◆ **Session Interval**—You can also specify an interval for additional transmissions by indicating a length of time and choosing minutes, hours, or days.
 - ◆ **Queue Threshold (percentage)**—Specify a percentage of the disk queue and system memory that must be in use to initiate the transmission process.
 - ◆ **Queue threshold (bytes)**—Specify the number of bytes that must be in the source queue and system memory to initiate the transmission process.
3. If no additional transmission options need to be set, establish the connection by clicking **Connect**.

Setting transmission stop criteria

Transmission stop criteria stops Storage Mirroring data transmission after a transmission start criteria has initiated the transmission. If a stop criteria has not been established, the transmission will end when the queue is empty.

1. Select the **Stop** option in the **Limit Type** box.



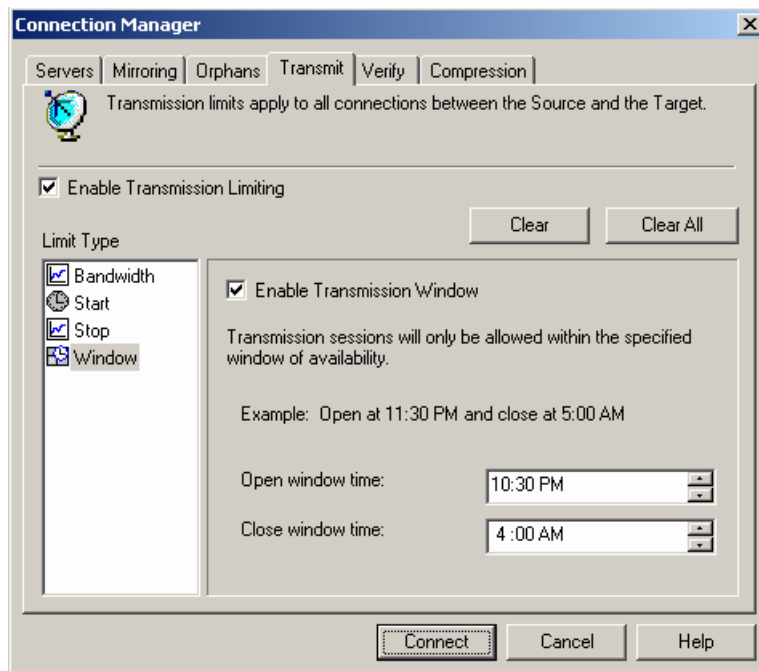
2. Define the stop options to stop Storage Mirroring transmissions by using either or both of the following options:
 - ◆ **Time Limit**—Specify the maximum length of time that Storage Mirroring can continue transmitting by indicating a length of time and choosing minutes, hours, or days.
 - ◆ **Byte Limit**—Specify the maximum number of bytes that can be sent before ending the Storage Mirroring transmission.
3. If no additional transmission options need to be set, establish the connection by clicking **Connect**.

Setting a transmission window

The transmission window establishes a period of availability for all Storage Mirroring transmissions.

NOTE: Setting a transmission window by itself is not sufficient to start a transmission. You still need to set a start criteria.

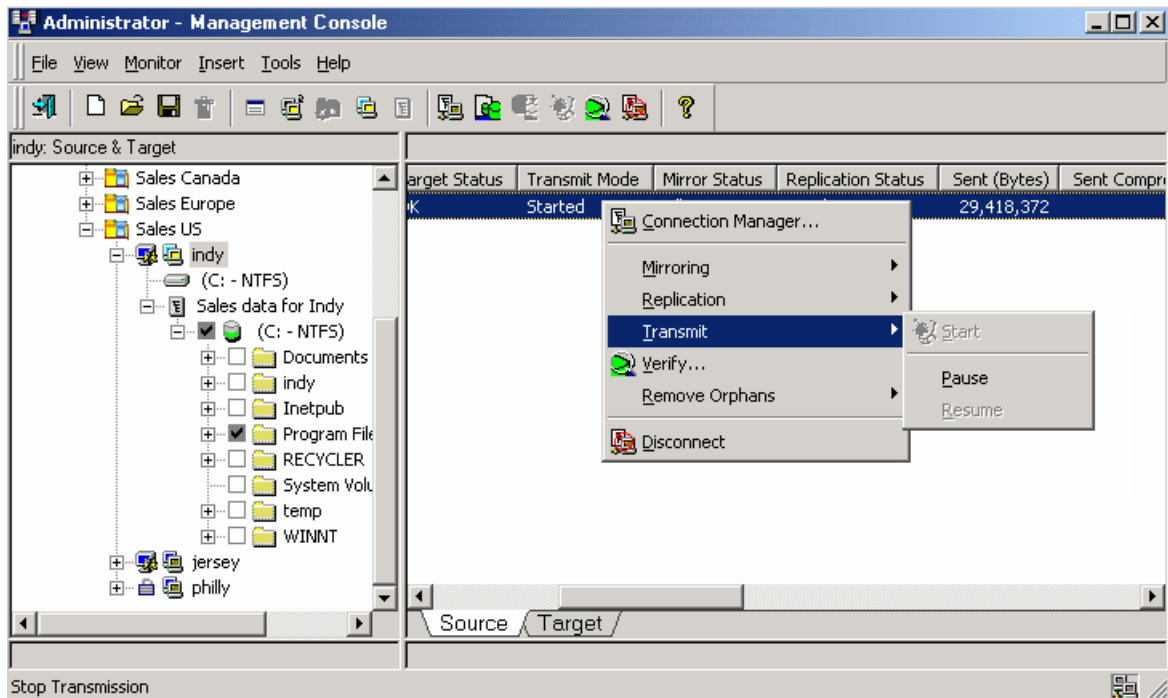
1. Select the **Window** option in the **Limit Type** box.



2. Mark the **Enable Transmission Window** check box to enable the transmission window feature.
3. Define the window of availability for Storage Mirroring transmission by setting the **Open window time** and **Close window time** settings. The time fields are formatted for any AM or PM time.
4. If no additional transmission options need to be set, establish the connection by clicking **Connect**.

Controlling transmission manually

To start, pause, or resume the transmission of data from the source to the target, right-click an established connection and select **Transmit** and the appropriate transmission control.



Setting Transmission Criteria Using the Text Client

The Text Client uses the `schedule` commands to establish scheduled transmission criteria. To control transmission manually using the Text Client, the `transmission` command can be used with the `stop`, `start`, `pause`, or `resume` switches.

Configuring transmission criteria before establishing a connection

After you have specified the source you are working with, establish scheduled transmission criteria using the following various `schedule` commands. After you have completed the transmission criteria, establish a source/target connection by following the steps in [Establishing a new connection](#) on page 6-16.

Clearing all transmission criteria

To clear all existing transmission options, use the `schedule clear` command.

Command	<code>SCHEDULE CLEAR</code>
Description	Clears the existing transmission schedule for the specified or current target machine
Syntax	<code>SCHEDULE <target_machine> CLEAR</code>
Options	target_machine —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>schedule jersey clear</code>

NOTE:	All transmission options are stored on the source machine until they are either cleared or the option is updated. At this time, there is not a command to display the existing transmission options.
--------------	--

Enabling and disabling transmission criteria

Enable transmission limiting when you want to apply any transmission options that have been configured. You can also disable the transmission options without losing your settings. Use either of the following two commands to disable or enable transmission limiting:

- ◆ Enable all existing transmission options by using the `schedule enable` command.

Command	<code>SCHEDULE ENABLE</code>
Description	Enables the transmission schedule
Syntax	<code>SCHEDULE <target_machine> ENABLE</code>
Options	target_machine —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>schedule jersey enable</code>

-
- ◆ Disable all existing transmission options by using the `schedule disable` command.

Command	<code>SCHEDULE DISABLE</code>
Description	Disables the transmission schedule without clearing the schedule data
Syntax	<code>SCHEDULE <target_machine> DISABLE</code>
Options	target_machine —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>schedule jersey disable</code>

Setting bandwidth limitations

Set the bandwidth by using the `limit bandwidth` command.

Command	<code>LIMIT BANDWIDTH</code>
Description	Sets bandwidth restrictions for transmitting data from the source to the target
Syntax	<code>LIMIT BANDWIDTH <bytes>, <seconds> TO <target_machine></code>
Options	<ul style="list-style-type: none">◆ bytes—Number of bytes to be transmitted◆ seconds—Maximum number of seconds to wait before transmitting again◆ target_machine—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>limit bandwidth 19300, 5 to jersey</code>
Notes	This command transmits in bursts, not bytes per seconds. The time identifies how long to wait before transmitting again. For example, if 5 seconds are specified and it only takes 2 seconds to send the specified bytes, Storage Mirroring will wait an additional 3 seconds before transmitting again.

Setting transmission start criteria

Specify any combination of the following options to start transmitting Storage Mirroring data to the target machine.

- ◆ **Specifying a Start Time**—Specify a start time to begin transmitting Storage Mirroring data using the `schedule start` command with the `starttime` switch.
- ◆ **Specifying a Start Time and Repeat Interval**—Specify a start time to begin transmitting Storage Mirroring data and a repeat interval for additional transmissions by using the `schedule start` command with the `starttime` and `every` switches.
- ◆ **Specifying Queue Usage**—Specify the queue usage criteria that must be in use to trigger Storage Mirroring data transmission. By using the `memlimit` switch, you can specify the percentage of the disk queue and system memory that must be in use to trigger transmission. By using the `queuesize` switch, you can specify the number of bytes that must be in the queue to trigger transmission.

Command `SCHEDULE START`

Description Sets criteria to start the transmission of data from the source to the target

Syntax `SCHEDULE <target_machine> START [STARTTIME = <mm/dd/yy> <hh:mm>]
[MEMLIMIT = <percent>] [QUEUESIZE = <bytes>] [EVERY <number> <time_units>]`

Options ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.

◆ **mm/dd/yy**—Date in month/day/year format indicating when the transmission will begin

◆ **hh:mm**—Time in hour:minute format using the 24 hour clock indicating when the transmission will begin

◆ **percent**—Any number between 0 and 100 indicating the percentage of the disk queue and system memory that must be in use to initiate the transmission process

◆ **bytes**—Number of bytes that must be in the source queue and system memory to initiate the transmission process

◆ **number**—Any number indicating how often the transmission process will be repeated

◆ **time_units**—Minutes (`min`), hours (`hr`), or days (`day`)

Examples ◆ `schedule jersey start starttime=3/11/04 03:30`
◆ `schedule philly start queuesize=400000`
◆ `schedule jersey start starttime=3/03/04 18:00, every 8 hr`
◆ `schedule philly start memlimit=25, queuesize=150000000`

Notes ◆ The start option **EVERY** cannot be used by itself and cannot be the first option in a string of options.

◆ If you use more than one start option, the transmission will begin when the first start option value is met. Additionally, each option after the first must be separated by a comma, as illustrated in the Examples.

Setting transmission stop criteria

Specify any of the following options to stop transmitting Storage Mirroring data to the target machine.

- ◆ **Specifying Transmission Duration**—Specify the maximum length of time that Storage Mirroring can continue transmitting by using the `schedule end` command with the `duration` switch.
- ◆ **Specifying the Maximum Number of Bytes**—Specify the maximum number of bytes that can be sent before ending the Storage Mirroring transmission by using the `schedule end` command with the `bytes` switch.

Command `SCHEDULE END`

Description Sets criteria to end the transmission of data from the source to the target

Syntax `SCHEDULE <target_machine> END [DURATION = <number> <time_units>][BYTES = <bytes>]`

Options ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.

◆ **number**—Any number indicating the length of time before the transmission ends

◆ **time_units**—Minutes (`min`), hours (`hr`), or days (`day`)

◆ **bytes**—Number of bytes transmitted before the transmission ends

Examples ◆ `schedule jersey end duration=3 hr`
◆ `schedule philly end bytes=1500000`

Notes If you use both of the end options, `duration` and `bytes`, the transmission will be stopped when the first end option value is met.

Setting a transmission window

Define the window of availability for Storage Mirroring transmission by using the `schedule window` command.

Command	<code>SCHEDULE WINDOW</code>
Description	Sets criteria to only allow transmissions during a certain period of time
Syntax	<code><u>SCHEDULE</u> <target_machine> WINDOW <hh:mm> TO <hh:mm></code>
Options	<ul style="list-style-type: none">♦ target_machine—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.♦ hh:mm—Time in hour:minute format using the 24 hour clock. The first time is when the transmission will begin and the second time is when the transmission will end.
Examples	<code>schedule jersey window 23:00 to 06:00</code>
Notes	Establishing a transmission window by itself is not sufficient to start a transmission. You will need to specify a start criteria.

Transmission schedule examples

This section shows examples of how the different `schedule` commands could be used together.

NOTE: The numbers and times were randomly selected for these examples. Be sure to use values that work for your environment.

- ◆ **Example 1**—In this example, transmission is set to begin on January 12, 2005, at 10:30 p.m. and to end after 6 hours. This schedule is also set to repeat every day:

```
schedule indy start starttime=1/12/05 22:30, every 1 day
schedule indy end duration=6 hr
schedule indy enable
```

If all the data is not transmitted within the 6-hour duration, the remaining data will remain in the queue and will be transmitted during the next scheduled transmission.

- ◆ **Example 2**—In this example, transmission is set to begin after the source queue contains 41943040 bytes (40 MB) of data, and transmission is set to end after 52428800 bytes (50 MB) of data have been sent from the source to the target:

```
schedule indy start queuesize=41943040
schedule indy end bytes=52428800
schedule indy enable
```

If there is data remaining in the source queue after the transmission ends, the data will be sent when the source queue again reaches 41943040 bytes (40 MB) of data.

- ◆ **Example 3**—In this example, transmission is set to begin after the source queue contains 52428800 bytes (50 MB) of data, and transmission is set to end after 62914560 bytes (60 MB) of data have been sent from the source to the target. However, transmission can only occur if the start criteria is met within the defined 6-hour window:

```
schedule indy start queuesize=52428800
schedule indy end bytes=62914560
schedule indy window 22:00 to 04:00
schedule indy enable
```

If additional data remains in the source queue after the transmission ends, that data will be sent when the start criteria is again met within the defined 6-hour window. However, if the start criteria is not met within the defined 6-hour window, data remains in the queue until the start criteria is met within the defined window.

NOTE: You do not have to specify the exact number of bytes in 1 MB. You can approximate. For example, the following is acceptable:

```
schedule indy start queuesize=5000000
schedule indy end bytes=6000000
```

At any time, transmission can be manually started, stopped, paused, or resumed regardless of scheduled transmission criteria.

Controlling transmission manually

To start, stop, pause, or resume the transmission of data from the source to the target, use the `transmission` command with the `start`, `stop`, `pause`, or `resume` switches.

- ◆ To start the transmission process, use the `transmission start` command.

Command	<code>TRANSMISSION START</code>
Description	Initiates the transmission process
Syntax	<code>TRANSMISSION START <target_machine></code>
Options	<code>target_machine</code> —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.

Examples `transmission start jersey`

- ◆ To stop the transmission process, use the `transmission stop` command.

Command	<code>TRANSMISSION STOP</code>
Description	Stops the transmission process
Syntax	<code>TRANSMISSION STOP <target_machine></code>
Options	<code>target_machine</code> —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.

Examples `transmission stop jersey`

- ◆ To pause the transmission process, use the `transmission pause` command.

Command	<code>TRANSMISSION PAUSE</code>
Description	Pauses the transmission pause
Syntax	<code>TRANSMISSION PAUSE <target_machine></code>
Options	<code>target_machine</code> —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.

Examples `transmission pause jersey`

- ◆ To resume the transmission process, use the `transmission resume` command.

Command	<code>TRANSMISSION RESUME</code>
Description	Resumes a paused transmission
Syntax	<code>TRANSMISSION RESUME <target_machine></code>
Options	<code>target_machine</code> —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.

Examples `transmission resume jersey`

Pausing Storage Mirroring Execution on the Target

Storage Mirroring servers running version 4.2 or later can pause the execution of operations on the target, even if there are no active connections to that target. For targets that do have active connections, any operations already in progress are completed, while any new incoming operations are queued until execution on the target is resumed, at which time the queued operations are processed.

Pausing through the Management Console

Use either of the following methods, to pause execution on the target:

- ◆ Right-click on the left pane of the Management Console, on a target server which you are logged in to, and select **Pause Target**.
- ◆ Right-click on the Target tab on the right pane of the Management Console, on a connection, and select **Pause Target**.

If you have an established connection, verify in the **Target Status** that the target is **Paused**. (This field is not updated until there is source/target activity.)

NOTE: The **Target Status** will identify **Pause Pending** while it completes the execution of operations that it was currently processing.

If you have multiple connections to the same target, all connections will be paused when that target is paused.

Pausing through the text clients

Pause the execution of operations on the target by using the `pausetarget` command on a target that you are logged in to.

Command	PAUSETARGET
Description	Allows you to pause the execution of Storage Mirroring operations on the target
Syntax	PAUSETARGET <target_machine> [FROM <source_machine>]
Options	<ul style="list-style-type: none">◆ target_machine—The name of the target machine where you want to pause execution of the Storage Mirroring operations◆ source_machine—The name of the source machine that is connected to the target
Examples	<code>pausetarget jersey</code>
Notes	<ul style="list-style-type: none">◆ You must be logged on to the target machine for this command to work.◆ If the target machine has not been identified using the <code>target</code> command, you must specify the <target_name> in the <code>pausetarget</code> command.

Verify that **(Paused)** is displayed after the target name on the Text Client screen.

Resuming through the Management Console

Use either of the following methods, to resume execution on the target:

- ◆ Right-click on the left pane of the Management Console, on a target server which you are logged in to, and select **Resume Target**.
- ◆ Right-click on a connection in the Target tab on the right pane of the Management Console and select **Resume Target**.

If you have an established connection, verify in the **Target Status** field that the target is resumed. (This field is not updated until there is source/target activity.)

NOTE: If you have multiple connections to the same target, all connections will be resumed when that target is resumed.

Resuming through the text clients

Resume the execution of operations on the target by using the `resumetarget` command on a target that you are logged in to.

Command	<code>RESUMETARGET</code>
Description	Allows you to resume the execution of Storage Mirroring operations on the target
Syntax	<code>RESUMETARGET <target_machine> [FROM <source_machine>]</code>
Options	<ul style="list-style-type: none">◆ target_machine—The name of the target machine where you want to resume execution of the Storage Mirroring operations◆ source_machine—The name of the source machine from where you want to resume execution of the Storage Mirroring operations
Examples	<code>resumetarget jersey</code>
Notes	<ul style="list-style-type: none">◆ You must be logged on to the target machine for this command to work.◆ If the target machine has not been identified using the <code>target</code> command, you must specify the <code><target_machine></code> in the <code>resumetarget</code> command.

Verify that (Paused) is no longer displayed after the target name on the Text Client screen.

Failover is a component of Storage Mirroring that allows a target to stand in for a failed source machine. The failover target assumes the network identity of the failed source. When the target assumes the identity of the source, user and application requests destined for the source machine or its IP address(es) are routed to the target.

When partnered with the Storage Mirroring data replication capabilities, failover routes user and application requests with minimal disruption and little or no data loss. In some cases, failover may be used without data replication to ensure high availability on a machine that only provides processing services, such as a web server.

NOTE: Failover is not supported in Replication Editions of Storage Mirroring.

The Failover Process

Storage Mirroring failover monitors the status of machines by tracking network requests and responses exchanged between a monitored source machine and the failover target. The time between requests and the number of allowable responses that can be missed combine to create a timeout period. When the source machine fails to respond before the timeout period has expired, Storage Mirroring determines that the source has failed. At this time, you will be prompted to initiate failover or it may occur automatically, if configured. In the event of failover, the target assumes or adds the identity of the failed source including machine name, IP address, and subnet mask.

Failover also sends updates to routers and other machines to update the IP to MAC address mapping. Network packets and applications destined for the failed IP address are routed to the target machine.

Depending on the type of client workstations, the timeout settings, and the applications in use, the clients may notice only a slight pause while the failover process occurs. If the failover timeout is set to a duration such as several minutes, clients may see an Abort or Retry message at their machine if they try to communicate with the source before the timeout has expired and the failover process has completed. For most 32-bit clients and network aware applications, reconnection is automatic. In the case of older client software, the workstation may need to be rebooted to reestablish a connection to the target, which is now acting as the source.

By incorporating user-defined failover scripts into the process, network administrators can automate many network and application events on the target machine, such as starting applications or system services, adding or removing IP addresses, or sending network messages to administrators.

Each Storage Mirroring target tracks which source machines it has failed over for. In the event that one of these machines comes back online on its own before the failback process has been completed, failover will not allow the source to mirror or replicate so that data integrity on the target is not compromised.

The Failback Process

The source machine problem(s) must be corrected while disconnected from the network to avoid a name or IP address conflict. After the problem(s) are corrected, the network administrator manually initiates failback. The source machine should not be reattached to the network until failback has completed. For Windows, this means that Storage Mirroring has completely removed the source's identity from the target.

Depending on the type of machine and data that Storage Mirroring is protecting, failback may need to be scheduled for an inactive period. If failover is being used in conjunction with Storage Mirroring replication or if a drive on the source was replaced, the data on the source may not be the most current information. It may be necessary to restore the most recent data from the target machine to the proper location on the source before initiating the failback process and bringing the source back online. (With Storage Mirroring, the restore can be accomplished by mirroring the changed data back to the source machine. See [Restoration](#) on page 12-1 for more information.)

Users may notice an interruption at their workstations during failback. This delay will occur between the completion of the failback process and the time needed to bring the source machine back online.

Like failover, network administrators can incorporate user-defined failback scripts into the process to automate many network and application events on the target machine, such as starting applications or system services, adding or removing IP addresses, or sending network messages to administrators.

Failover Clients

- ◆ **Text Client**—Failover processing is integrated in DTCL allowing you to:
 - ◆ Configure both monitoring and failover options from any of the Storage Mirroring text clients
 - ◆ Configure failover for any operating system
 - ◆ Intiate failover and failback
- ◆ **Failover Control Center**—The Failover Control Center is a standalone GUI client allowing you to:
 - ◆ Configure both monitoring options and failover options
 - ◆ Display the status of all monitored machines and their IP addresses
 - ◆ Initiate failover and failback


NOTE: If you are using Active Directory, you must use the Failover Control Center to configure your failover settings.

Failover Terminology


- ◆ **MAC Addresses**—The MAC (media access control) address represents the physical station or hardware address of the NIC (network interface card). These addresses are used to identify computers on a network. MAC addresses are unique and are coded into every network adapter by its manufacturer.
- ◆ **Missed Packets**—The missed packets is the configurable setting that specifies how many monitor replies can be missed before assuming the source machine has failed.
- ◆ **Monitored IP Addresses**—Monitored IP addresses are the IP addresses on the source selected to be monitored by a Storage Mirroring target.
- ◆ **Monitor Interval**—The monitor interval is the configurable setting that specifies how often the monitor request is sent to the source machine.
- ◆ **Monitor Reply**—A monitor reply is the acknowledgment that the source sends to the target indicating that it is still online.
- ◆ **Monitor Request**—A monitor request is an inquiry sent from the target to the source to determine if the source is still active.
- ◆ **Shares**—A share is any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users.
- ◆ **Timeout**—The failover timeout is the amount of time before failover begins. The time is calculated by multiplying the values of the monitor interval and missed packets settings. This time is displayed on the Failover Control Center.
- ◆ **Unmonitored IP Addresses**—On a machine with multiple IP addresses, Storage Mirroring will not monitor the IP addresses that are not selected, but the target can assume these addresses during failover, if configured.

How Failover Works

1. The target machine monitors the source machine and waits for a failure.




2. A source failure is detected and the failover process begins. (For more information on failure detection, see [How Failure Detection Works](#) on page 11-9.)
3. If specified, a pre-failover script executes on the target. This script is user-defined and optional. Scripts may be used to stop services on the target that will not be needed or may be in conflict after the target stands in for the failed source. Scripts can also send network messages notifying administrators that failover is about to occur.



NOTE: Until the failover process has completed executing, user and application requests destined for the source machine will not be delivered.

4. Depending on the source and target machine operating system, one of the following failover events will occur:
 - ◆ **Windows 200x**—The failed source's machine name, IP address, and subnet mask are added to the target's identity.
 - ◆ **Windows NT 4.0**—The failed source's machine name, IP address, and subnet mask are added to the target's identity. Although there is an option to replace the target's identity with that of the source, the add option is recommended to ensure that there are no IP address conflicts upon failback.



-
5. Source shares are created on the target allowing shared resources on the source to be accessible on the target machine. Any drives shares that existed on the target prior to failover are stored so that they can be reestablished after failback.

NOTE: When replicating data using Storage Mirroring, only shares that are within the Storage Mirroring replication set will be created on the target machine during the failover process. Shares outside of the replication set will not be created on the target.

6. Failover sends updates to routers and other machines with IP to MAC address mappings allowing clients to seamlessly attach to the target machine. All user and application requests destined for the source machine are routed to the target machine.
7. If specified, a post-failover script executes on the target. This script is user-defined and optional. This may be used to start services that will be needed after the target assumes the identity of the failed source. Scripts can also send network messages notifying administrators that failover is complete.
8. Failover is complete.

How Failback Works

1. The network administrator disconnects the failed source machine from the network and repairs it.
2. The network administrator manually initiates failback.
3. If specified, a pre-failback script executes on the target. This script is user-defined and optional. This may be used to stop services on the target that will not be needed or may be in conflict after the target reinstates its own identity. Scripts can also send network messages notifying administrators that failback is about to occur.
4. The source's identity is removed from the target or the mapped aliases and drives are removed.

NOTE: Users who were using the name or address of the source machine were actually attached to the target machine and will lose their connection to the server at this point.


5. For Windows, shares created on the target during failover are deleted. Any original shares that existed are restored.
6. Failback sends updates to routers and other machines with IP to MAC address mappings. This mapping is now the original mapping before failover occurred. At this time, users may see a pause at their workstations while the failback process completes.
7. If specified, a post-failback script executes on the target. This script is user-defined and optional. This may be used to start services that will be needed after the target reinstates its own identity. Scripts can also send network messages notifying administrators that failback is complete.
8. Failback is complete and you are given the choice to continue monitoring the source machine.
9. The original source machine can be reattached to the network.

WARNING: After the source machine is back online, if you need to restore the data on the original machine by mirroring data from the target's copy of the replication set, do not let users logon to the source until the restore process is complete.

10. If specified, a post-failback script executes on the source. This script is user-defined and optional.
11. Failback is complete.

How Failure Detection Works

1. The target sends a monitor request, which is an ICMP ping or a UDP heartbeat, to each monitored IP address at a user-defined interval. A monitor reply is sent from the source back to the target.
2. When the user-defined number of missed packets is met, the address is considered failed.



Client Connectivity

Depending on the type of client workstations, the timeout settings, and the applications in use, the clients may notice only a slight pause while the failover process occurs. If the failover timeout is set to a duration such as several minutes, clients may see an Abort or Retry message at their machine if they try to communicate with the source before the timeout has expired and the failover process has completed. For most 32-bit clients and network aware applications, reconnection is automatic. In the case of older client software, the workstation may need to be rebooted to reestablish a connection to the target, which is now acting as the source.

Operating System Differences

The source and target machines should be similarly configured (operating system, applications, service packs, patches, and so on) in order for the applications and services running on the source to be available on the target after failover. Since the Windows registry should not be replicated between machines, required applications should be preinstalled on the target and left idle so they do not open or modify replicated files. The pre-failover and post-failover scripts will start required applications. The following pages detail some of the operating system differences of failover.

Windows 200x

- ◆ Storage Mirroring failover can monitor an unlimited number of source machines and can assume the identity of multiple sources.
- ◆ During the failover process the target maintains its original identity and adds the name(s) and IP address(es) of one or more failed source machines.
- ◆ Global unique IDs (GUID) used in Windows 200x are mirrored and replicated to the target machine, stored in an attribute database, and only set if failover occurs.
- ◆ If the identity of the source is added to the target during failover, no services are stopped or restarted.
- ◆ The following table describes the source and target configurations of Windows 200x machines that are supported by Storage Mirroring.

		Target	
		Domain Controller	Member
Source	Domain Controller	X ^a	X
	Member		X

- a. If you are using Active Directory, additional steps are required in your failover and failback scripts to accommodate Active Directory security.

Windows NT 4.0

- ◆ Storage Mirroring failover can monitor an unlimited number of source machines (limited only by the number of IP address placeholders that exist) and can assume the identity of multiple sources.
- ◆ During the failover process, by default, the target can maintain its original identity and add the name(s) and IP address(es) of one or more failed source machines.
- ◆ The target can optionally be configured to replace its original identity with that of the source machine. When replace is selected, the target can only assume the identity of one source at a time. It is, however, possible to fail over for other sources by adding their identity. Replacing is available for applications that rely on the primary machine name and are sensitive to machine name changes. It is possible to replace the target identity with one source and then add the identity of other source machines.
- ◆ If the identity of the source is added to the target during failover, no services are stopped or restarted. If the identity of the target is replaced by the source, the Computer Browser, Net Logon, and Server services are stopped and restarted during both failover and failback. Other services may be stopped or restarted as needed using failover and failback scripts.
- ◆ If the identity of the source replaces the target during failover, the target machine must be a PDC/BDC or the target will lose authentication with the domain.
- ◆ The following table describes the source and target configurations of Windows NT 4.0 machines that are supported by Storage Mirroring.

		Target		
		Primary Domain Controller	Backup Domain Controller	Member
Source	Primary Domain Controller		X ^a	X ^b
	Backup Domain Controller	X	X	X ^c
	Member	X	X	X

- a. After failover, the target will have the functionality of a BDC, not a PDC, but could be promoted to a PDC.
- b. After failover, the target will have the functionality of a member, not a PDC. It is necessary to promote a BDC so that users/groups can be added, deleted, or modified. The user will not be able to log into the domain unless there is a BDC.
- c. After failover, the target will have the functionality of a member, not a BDC.


Failover Hardware Configurations

Failover can be configured to stand in for one or more IP addresses associated with different NICs on the source. Each IP address can be added to a specific target NIC making NIC configuration very flexible. Following are some sample configurations.

NOTE: If a failover target is monitoring a source for failure and a new NIC is added to the target server, you will need to reinstall Storage Mirroring to properly setup the new NIC for failover placeholders. An uninstall is not necessary; you can reinstall by choosing the Repair or Upgrade option during the installation. For detailed information on repairing or upgrading Storage Mirroring, see the Storage Mirroring *Getting Started* guide.

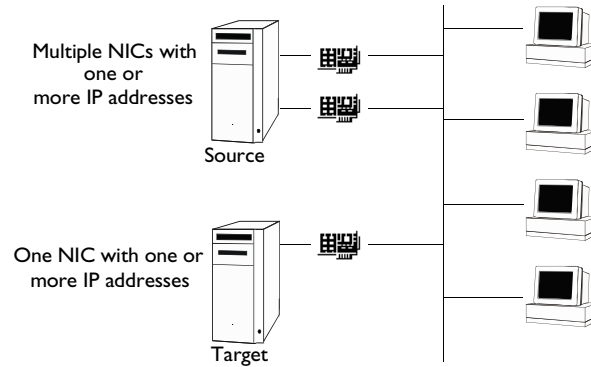
Single NIC on the source and target

A single NIC on the source may have one or more IP addresses assigned to it. If that source or the NIC fails, all traffic from the source is directed to the target. If there are multiple IP addresses on one source or if there are multiple sources, the target has to assume traffic from all of the addresses.



Multiple NICs on the source and single NIC on the target

Multiple NICs on the source can increase the impact on the target in the event of a single NIC failure. Using failover monitoring options, network administrators can specify how to monitor addresses on separate NICs and trigger failover if any or all monitored IP addresses have failed. In the event of an entire machine failure, the target has to assume the traffic from both NICs and their assigned IP addresses, in addition to the traffic it is already processing.




WARNING: With multiple source NICs or IP addresses you must carefully configure your monitoring and failover settings to avoid causing IP address conflicts or compromising data integrity. To eliminate any potential for data integrity problems or IP address conflicts, it is suggested that you monitor all IP addresses on the source machine, and trigger failover when all IP addresses fail, when it is possible. For more information, see [Failover Options](#) on page 11-15 for information on **Failover Trigger** and **Items to Failover**.

Multiple NICs on the source and target

With any source configuration, additional NICs on the target increases flexibility and control. Secondary target NICs can assume the traffic from a failed source NIC while normal target traffic can continue to use the primary target NIC. Multiple sources can be directed to specific NICs.

For example, a single target containing three NICs monitors three source machines located on different subnets. Each source has one NIC attached to a different subnet. Each source's IP address is configured to be assumed by the target NIC on the appropriate subnet.



NOTE: When using multiple subnets, the location of the source and target can be critical. If the source and target are located on different subnets, users cannot access the source machine name and IP address on the target's subnet after failover because of routing conflicts. To avoid this situation, keep the source and target on the same subnet or use a virtual private network.

Failover Options

In order for failover to stand in for a failed machine, you must configure the monitoring and failover options. These options are the same regardless of the configuration method (Failover Control Center or Text Client) that you use.

- ◆ **Source Machine**—The machine that is monitored for failure.
- ◆ **IP Addresses**—The machine that is monitored for failure may have one or more IP addresses. When multiple IP addresses are available, you have the option of monitoring some or all of the addresses as well as failing over for some or all of the addresses.
- ◆ **Target Machine**—The machine that will be assuming the identity of the source machine in the event that the source machine fails.
- ◆ **Target NIC**—The network card on the target machine that will receive the traffic from the failed source machine.
- ◆ **Monitoring Method**—You can choose to monitor the source by using ICMP pings or UDP heartbeats. You can also choose to disable monitoring, in which case you will be responsible for identifying when a failure has occurred and initiating failover manually.
- ◆ **Monitor Interval and Missed Packets**—The monitor interval specifies how often the monitor request is sent to the source machine. The missed packets specifies how many monitor replies can be missed before assuming the source machine has failed. To achieve shorter delays before failover, use lower monitor interval and missed packets values. This may be necessary for IP addresses on machines, such as a web server or order processing database, which must remain available and responsive at all times. Lower values should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, lower values can lead to premature failover. To achieve longer delays before failover, choose higher monitor interval and missed packet values. This may be necessary for IP addresses on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.
- ◆ **Failover Trigger**—When multiple IP addresses on a source machine are being monitored, failover can begin after one monitored address fails or it can be delayed until all of the monitored addresses fail. Consider that if there are multiple redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail. You should also consider that if each IP address is on a different subnet, you should trigger failover after one fails.
- ◆ **Target Identity after Failover**—You have the option of adding the source's identity to the target machine or replacing the target identity with that of the source.
- ◆ **Failover Execution**—When multiple IP addresses on a source machine are available (whether they are all monitored or not), failover can occur for just the monitored IP addresses or for all of the machine's IP addresses. The same considerations specified in **Failover Trigger** should be kept in mind when configuring failover execution. The execution is dependent on the machine configuration when multiple NICs are in use. You can also specify if the server name and any source shares will be failed over. If you specify the server name to be failed over, first Storage Mirroring checks the hosts file and uses the first name there. If there is no hosts file, Storage Mirroring uses the first name in DNS. (Although, the first name in DNS may not always be the same each time the DNS server is rebooted.) Lastly, if there is no DNS server, Storage Mirroring uses the Failover Control Center monitor name.
- ◆ **Share Mapping**—You have the option of failing over shares so that they are available on the target machine after failover. If share mapping is not used, users will not be able to access shared data.

NOTE: Only standard file system shares are created on the target during failover. If you are using MSCS cluster shares created through the MSCS Cluster Administrator, those shares will need to be manually created on the target machine.

If you are configuring failover so that drive shares are created on the target during failover but your source and target machines do not have the same drive letters, you must use the **All-to-One** selection on the Servers tab of the Connection Manager or the `map base` command in the Text Clients when establishing your connection. Otherwise, the shares will not be created on the target during failover.

If a share is created on Windows NT 4 or Windows 2000 with the default full access permissions (without an ACL) and then failed over, the permissions given to the target will be dependent on the target's operating system. On a Windows 2003 target, the default permission will only be read only permission.

If you are failing over multiple source machines and they contain shares with identical names, the shares from the second failed source will replace the shares from the first failed source. For example, if two servers both have a share called `\\users`, and one machine fails, `\\users` from that first failed server will be available on the target after failover. If the second server fails, the `\\users` share will failover and replace the `\\users` share currently on the target (from the first failed machine). Use unique share names across all servers to avoid this situation.

-
- ◆ **Manual Intervention**—Manual intervention allows you to control when failover occurs. When a failure occurs, a prompt appears in the Failover Control Center and waits for you to manually initiate the failover process. You can disable manual intervention if you want failover to occur automatically.
 - ◆ **Scripts**—The failover and failback scripts for the target are stored on the target machine, but are unique for each source machine. The failback script for the source is stored on the source machine. Scripts may contain any valid Windows command, executable, or batch file. To see samples of failover and failback scripts, see [Application Failover](#) on page B-1.

NOTE: Failover scripts will run but will not display on the screen if the Storage Mirroring service is not set to interact with the desktop. Enable this option through the Services applet.

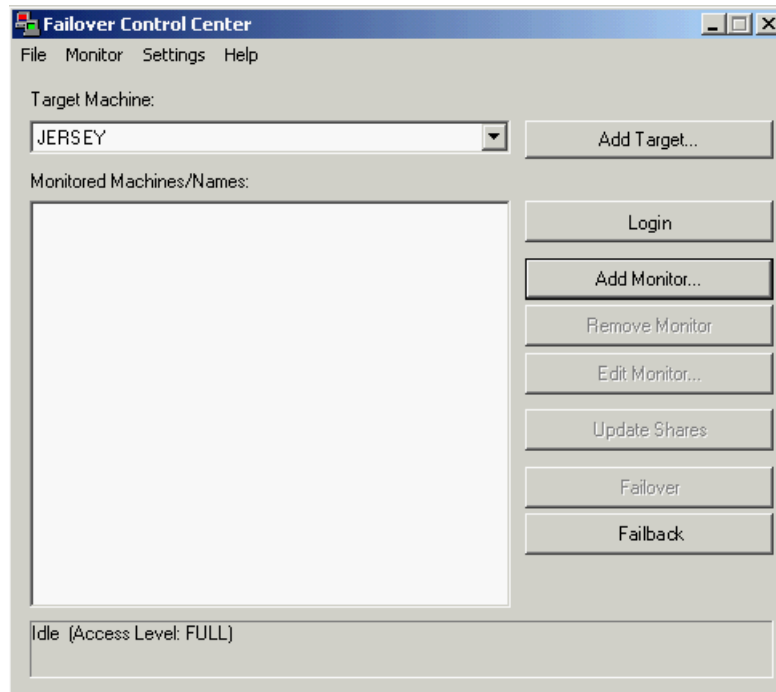
- ◆ **Active Directory Account**—If you are using Active Directory, you will need to use the Failover Control Center to identify a user and the associated password that has update privileges within Active Directory. This allows SPNs to be created and deleted during failover.

Managing Failover Through the Failover Control Center

Through the Failover Control Center, you can configure and monitor the failover process. The failover settings determine what users experience during a failover. In addition, the Failover Control Center can be used to initiate the failback process and to offer additional options.

Configuring monitoring and failover


1. The Failover Control Center can be started from within the Management Console or from the Windows desktop.
 - ◆ From the Management Console, select **Tools, Failover Control Center**.
 - ◆ From the Windows desktop, select **Start, Programs, Storage Mirroring, Failover Control Center**.



NOTE: You may need to log in to a target machine using your Storage Mirroring administrator username and password.


For detailed information on the Failover Control Center display, see [Understanding the Display](#) on page 4-3.

2. Select a failover target from the **Target Machine** list box. If the target you need is not listed, click **Add Target** and manually enter a name or IP address (with or without a port number). You can also select the **Browse** button to search for a target machine name.



3. Click **OK** to select the target machine and return to the Failover Control Center main window.
4. Select a source machine to monitor by clicking **Add Monitor**. The Insert Source Machine dialog box appears.

- Specify your source machine by any of the following methods:

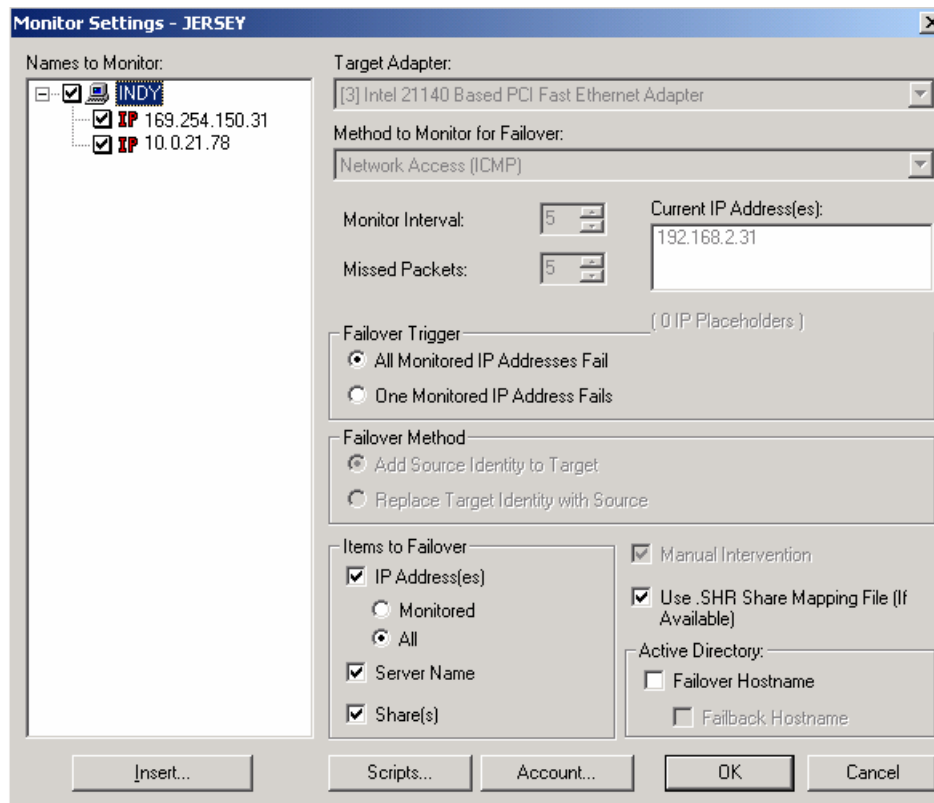


- Type the name of the machine that you want to monitor in **Machine Name(s)** and click **OK**. You can enter multiple names if you want to monitor more than one machine by separating the names with a space.
 - Click **Custom**, specify the IP address (with or without a port number) of the machine you want to monitor, click **OK**.
 - Click **Browse** to search for a machine. Select a domain from the list box at the top of the Select Machine dialog box to list the available machines for that domain. After highlighting a machine to be monitored (or selecting multiple machines with the Ctrl or Shift keys), click **OK**.
- The Monitor Settings dialog box appears.

NOTE: If you want to edit the monitor settings for a source machine that is currently being monitored, highlight that source machine on the **Monitored Machines** tree on the main Failover Control Center screen and click **Edit**. You will not get the Insert Source Machine dialog box, but will immediately open the Monitor Settings dialog box.

If you want to discontinue monitoring a source machine, highlight that machine on the **Monitored Machines** tree on the main Failover Control Center screen and click **Remove Monitor**. No additional dialog boxes will open.

Adding or editing a source monitor is completed through the Monitor Settings dialog box. This dialog box is used to configure all of the monitor settings.



The Monitor Settings dialog box is divided into five operational areas:

- ◆ Source Machines and IP Addresses—**Names to Monitor**
- ◆ Target Machine Information—**Target Adapter**, **Current IP Address(es)**, and **Method to Monitor for Failover**
- ◆ Monitor Requests and Responses—**Monitor Interval** and **Missed Packets**
- ◆ Failover Processing—**Failover Trigger**, **Failover Method**, and **Items to Failover**
- ◆ Miscellaneous Source Machine Options—**Manual Intervention**, **Use.SHR Share Mapping File**, and **Active Directory**

NOTE: The target machine is identified in the Monitor Settings dialog box title bar. All changes made in this dialog box are for the machine referenced in the title bar. If the wrong target is selected, close the Monitor Settings dialog box and select a different target from the Failover Control Center window.

7. Select an IP address or source to be monitored by marking the check box to the left of an item in the **Names to Monitor** tree.
8. Highlight one IP address in the **Names to Monitor** tree and select a **Target Adapter** that will assume that IP address during failover. Repeat this process for each IP address that is being monitored.

NOTE: **Current IP Addresses** displays the IP address(es) currently assigned to the selected target adapter.

9. Specify the **Method to Monitor for Failover**. You can choose to monitor the source by using ICMP pings or UDP heartbeats. You can also choose to disable monitoring, in which case you will be responsible for identifying when a failure has occurred and initiating failover manually. If the server does not support non-ICMP monitoring methods, this option will not be configurable.

NOTE: If you choose to use UDP heartbeats, your source and target must be using the same Storage Mirroring port settings. If you choose to use ICMP pings, the port settings can be different.

10. Specify the **Monitor Interval**. This setting identifies the number of seconds between the monitor requests sent from the target to the source to determine if the source machine is online. This option is not configurable if your **Method to Monitor for Failover** is set to **No Monitoring**.
11. Specify the **Missed Packets**. This setting is the number of monitor replies sent from the source to the target that can be missed before assuming the source machine has failed. This option is not configurable if your **Method to Monitor for Failover** is set to **No Monitoring**.
12. Repeat steps 8-11 to configure monitor settings for other source machine IP addresses.
13. If you are monitoring multiple IP addresses, highlight the machine name and specify the **Failover Trigger**.
 - ◆ **All Monitored IP Addresses Fail**—Failover begins when all monitored IP addresses fail.
 - ◆ **One Monitored IP Address Fails**—Failover begins when any of the monitored IP addresses fails.
14. Highlight the machine name and specify the **Failover Method** option by specifying how the source will impact the target machine:
 - ◆ **Add Source Identity to Target**—The source machine's identity is added to the target machine's identity.
 - ◆ **Replace Target Identity with Source**—The source machine's identity replaces the target machine's identity.

NOTE: If the **Replace Target Identity with Source** check box is selected and Storage Mirroring is replicating data from multiple source machines to a single target and one of the sources fail, the Storage Mirroring failover process automatically notifies the other Storage Mirroring sources of the new identity of the target so that replication can continue.

-
15. If you are monitoring multiple IP addresses, highlight the machine name and specify your **Items to Failover** options.

◆ **IP Addresses**

- ◆ **Monitored IP Addresses**—Failover is only performed on the monitored IP addresses.
- ◆ **All Addresses**—Failover is performed on the monitored and unmonitored IP addresses.

NOTE: When multiple IP addresses are monitored, there is the potential for IP address conflicts during failover. Conflict may occur when the number of IP addresses that trigger failover is less than the number of IP addresses that are assumed by the target during failover.

For example, suppose a source machine has three IP addresses and two of these are monitored. If a failover occurs due to a machine failure, there is not an address conflict because all of the IP addresses have failed and no longer exist. If the failure only occurs on one of the monitored addresses, the other two IP addresses (monitored and unmonitored) are still affected. The **Items to Failover** option forces these addresses to be assumed by the target as well. These addresses then exist on both the source and the target. This occurs because the **Items to Failover** option allows the selected (monitored) IP addresses or all of the machine addresses to be failed over, but it does not allow the individual failed IP address to be failed over. Therefore, when a source machine has fewer IP addresses that trigger failover than IP addresses that will be assumed, there is a risk of an IP address conflict.

The delay before the failover process begins can be used to determine the cause of the failover, thus possibly avoiding IP address conflicts.

- ◆ **Server Name**—Failover is performed on the server name.
- ◆ **Shares**—Failover is performed on shares.

NOTE: Share information is automatically updated on the target once an hour. If you need to manually update the share information, click **Update Shares** on the main Failover Control Center window after the connection has been established.

If a share is created on Windows NT 4 or Windows 2000 with the default full access permissions (without an ACL) and then failed over, the permissions given to the target will be dependent on the target's operating system. On a Windows 2003 target, the default permission will only be read only permission.

16. Disable **Manual Intervention** only if you want failover to occur immediately when a failure occurs. This option is not configurable if the **Method to Monitor for Failover** for monitored IP address is set to **No Monitoring** in step 9.
17. If the **Shares** selection under **Items to Failover** is selected, verify that the **Use .SHR Share Mapping File** check box is selected if you would like to use the Storage Mirroring share mapping file to create shares on the target machine during failover. If the **Use .SHR Share Mapping File** check box is not selected, shares will be created using the information gathered when the machine was selected as a source machine to be monitored.

NOTE: If the **Shares** selection under **Items to Failover** is not selected, shares will not be failed over to the target regardless of the **Use .SHR Share Mapping File** selection.

Share information is automatically updated on the target once an hour. If you need to manually update the share information regardless of the **Use .SHR Share Mapping File** selection, click **Update Shares** on the main Failover Control Center window after the monitor has been created.

Automatic share failover only occurs for Windows shares. Other shares, such as NFS shares or Macintosh Volumes, must be configured for failover through the failover scripts. For example scripts, see [Application Failover](#) on page B-1.

18. By default, **Failover Hostname** is disabled. This option automatically removes the host SPN (Service Principle Name) from Active Directory on the source and adds it to Active Directory on the target. If you are using Active Directory, enable this option or you may experience problems with failover.
19. **Failback Hostname** returns the host SPN on the source and target back to their original settings on failback. If you are using Active Directory, enable this option or you may experience problems with failback.
20. If you are using Active Directory, click **Account** and identify a user and the associated password that has update privileges within Active Directory. This allows SPNs to be created and deleted during failover and failback. The username must be in the format domain\user or user@domain. Click **OK** to return to the Monitor Settings dialog box.

NOTE: If you are using Active Directory, your account password cannot be blank.

-
21. If you are using any scripts, click **Scripts** and enter the path and filename for each script type.

The 'Scripts' dialog box is shown with the following details:

- Target Section:**
 - Pre-Failover:** c:\scripts\preover.bat (with browse button)
 - ☒ Delay failover until script completes
 - Post-Failover:** c:\scripts\postover.bat (with browse button)
 - Pre-Failback:** c:\scripts\preback.bat (with browse button)
 - ☒ Delay failback until script completes
 - Post-Failback:** c:\scripts\postback.bat (with browse button)
- Source Section:**
 - Post-Failback:** c:\scripts\postback_src.bat (with browse button)
- ☐ Make these scripts defaults for future monitor sessions
- Buttons:** OK, Cancel


NOTE: If you click **Browse** next to the **Source Post-Failback** script, the directory structure for the target machine appears by default. From this point, you can access your network so that you can see the directory structure for the source machine, and then select the file you want.

22. If you want to delay the failover or failback processes until the associated script has completed, mark the appropriate check box.
23. If you want the same scripts to be used as the default for future monitor sessions, mark the appropriate check box.
24. Click **OK** to return to the Monitor Settings dialog box and click **OK** again to save your monitor settings and begin monitoring for a failure.

Monitoring failover

Now that replication and failover monitoring are configured and started, you will need to know if and when there is a problem. Since it can be essential to quickly know the status of your machines, Storage Mirroring offers various methods for monitoring the status of failover. When the Failover Control Center is running, you will see four visual indicators:

- ◆ The Failover Control Center Time to Fail counter
- ◆ The Failover Control Center status bar located at the bottom of the window
- ◆ The Failover Control Center colored bullets to the left of each IP address and source machine
- ◆ The Windows desktop icon tray containing a failover icon



NOTE: You can minimize the Failover Control Center and, although it will not appear in your Windows taskbar, it will still be active and the failover icon will still appear in the desktop icon tray.

The Failover Control Center does not have to be running for failover to occur.

The following table identifies how the visual indicators change as the status of failover changes.

	Time to Fail Countdown	Status Bar	Colored Bullets	Desktop Icon Tray
Source is Online	The Time to Fail counter is counting down and resetting each time a heartbeat is received from the source machine.	The status bar indicates that the target machine is monitoring the source machine.	The bullets are green. ^a	The Windows desktop icon tray contains a failover icon with red and green computers.
Source Fails and Failover is Initiated	The Time to Fail countdown value is 0.	The status bar displays the source machine and IP address currently being assumed by the target.	The bullets are red.	The Windows desktop icon tray contains a failover icon with red and green computers.
Failover is Complete	The Time to Fail counter is replaced with the "Failed Over" message.	The status bar indicates that monitoring has continued.	The bullets are red.	The Windows desktop icon tray contains a failover icon with a red computer.

a. When the **Time to Fail** value has decreased by 25% of the entire timeout period, the bullet changes from green to yellow, indicating that the target has not received a response from the source. The yellow bullet is a caution signal. If a response from the source is received, the countdown resets and the bullets change back to green. If the countown reaches zero without the target receiving a response from the source, failover begins.

Initiating failback

When failover occurs, a source machine has failed. The steps below must be completed in order to complete failback without IP address or name conflicts.

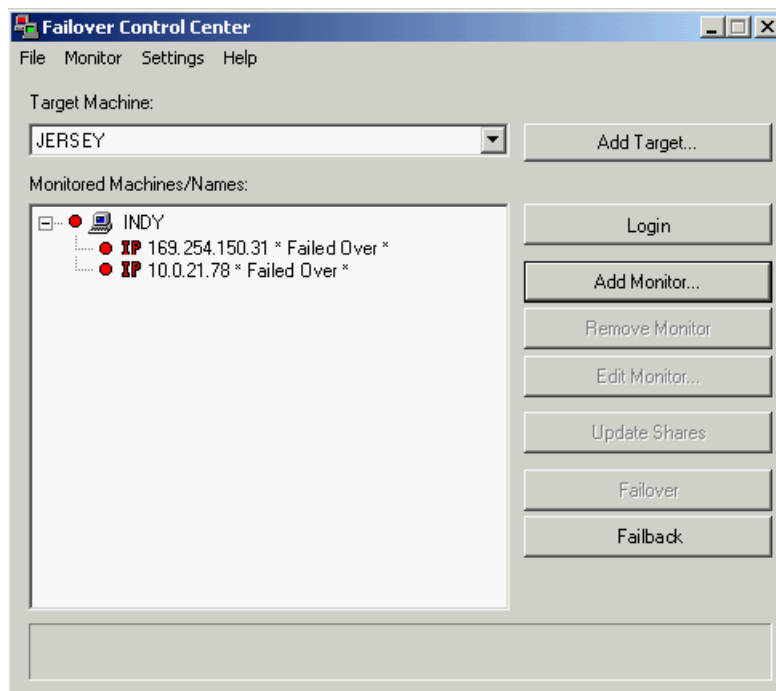
1. Verify that your source machine is not connected to the network. If it is, disconnect it.
2. Resolve the source machine problem that caused the failure.

WARNING: Do not connect the source machine to the network at this time.


3. In the Failover Control Center, select the target machine that is currently standing in for the failed source.

NOTE: If your target machine has been replaced during failover because of the configuration you selected (perform failover by replacing target identity with source), do not select your target from the **Target Machine** list. Since the target has been replaced by the source, it no longer exists and the Failover Control Center will retry for several minutes (possibly up to five minutes), until it finally fails in locating the target. After a replace failover, select your source machine in the **Target Machine** list, which will be monitoring itself in a failed over state, and perform failback. After failback, the target will be available again.

4. Select the failed source and click **Failback**. If you specified a pre-failback script in your failover configuration, that script will be executed at this time.



5. You will be prompted to determine if you want to continue monitoring the source. Bring the source online and select **Continue** or **Stop** to indicate if you want to continue monitoring the source.



After you have selected whether or not to continue monitoring the source machine, the source post-failback script, if configured, will be started.

NOTE: At this time, the source machine must be online and Storage Mirroring must be running to ensure that the source post-failback script can be started. If the source has not completed its boot process, the command to start the script may be lost and the script will not be initiated.

Failback is now complete. If you need to restore newer files from the target to the source, see [Restoring Storage Mirroring Data Through the Management Console](#) on page I2-5.

Failover IP address placeholders

For Windows NT 4.0, IP address placeholders exist so that during the failover process, the source's IP address can be assumed by the target without requiring a target machine reboot. The number of placeholders defined must be the maximum number of IP addresses for which the target adapter will need to respond. If all of the defined placeholders have been replaced due to failovers, the target cannot assume or respond for any additional IP addresses.

For example, suppose the following:

- 10 placeholders defined
- 3 machines being monitored
- 2 network adapters on each machine
- 2 IP addresses assigned to each adapter


Therefore:

2 IP addresses X 2 network adapters X 3 machines = 12 total IP addresses

If Storage Mirroring is set to monitor each IP address, on each adapter, on each machine, and all three machines failed, there are not enough placeholders available for the target to respond for every address. The last two IP addresses to fail do not get assumed by the target. Placeholders may have been created during the Storage Mirroring installation depending on the options you chose. The maximum number of placeholders allowed is 99.

Updating the number of IP address placeholders

1. From the Failover Control Center window, select **Settings, Placeholders**.



2. Modify the setting if necessary. Click **Set** to save the new configuration and reboot the machine, or **Close** to exit without saving the new setting.

NOTE: When the **Set** button is selected, a message box prompts to reboot the machine. A reboot is necessary to reinitialize the target adapter.

- ◆ Select **Yes** to set the new placeholder value and reboot the machine.
- ◆ Select **No** to set the new placeholder value, but a manual reboot will still be required to use this machine as a failover target.
- ◆ Select **Cancel** to not reboot the machine. The new placeholder value will not be set.


The number of placeholders is set independently for each NIC on a Windows NT 4.0 machine. For example, if you change the number of placeholders to 7 for one NIC, all other NICs on that machine will automatically be changed to 7 placeholders.

Communication ports

The Storage Mirroring service and the Failover Control Center use two ports for various Storage Mirroring communications.

- ◆ **Service Transmit Port**—The Failover Control Center uses this port to send commands to Storage Mirroring servers.
- ◆ **Heartbeat Listen Port**—The Failover Control Center uses this port to listen for heartbeats transmitted from Storage Mirroring servers.

To view or modify the port settings in the Failover Control Center, select **Settings, Communications**.



NOTE: Since the Management Console and the Failover Control Center share port settings, if the Storage Mirroring service is restarted, changes to the port settings that were made in the Failover Control Center are reflected in the Management Console.

If the Storage Mirroring service is not restarted, the changes will occur in the Failover Control Center only.

Manual share updates

Share information can be manually updated from the Failover Control Center window. To manually update the share information, highlight a source machine in the **Monitored Machines** tree and click the **Update Shares** button.


Failover Control Center window refresh rate

The failover client periodically requests information from the source and target. Depending on the type of information, the request may be a machine-specific request, like obtaining the **Time to Fail** status from a target, or may be a general request, like determining which machines are running Storage Mirroring.

The rate at which these requests are made can be modified through the Failover Control Center. Select **Settings, Refresh Rate**. The default update interval is one second.

A lower refresh rate value updates the information in the Failover Control Center window's **Monitored Machines** tree more often, but also generates more network traffic and higher utilization on the client and target machines.

A higher refresh rate value updates the information less frequently, but minimizes the network traffic.




Testing failover

The failover process, including pre-failover and post-failover script processing, can be tested at any time. To force unavailability, disconnect the network cable from a monitored machine, wait for the **Time to Fail** counter to decrease to zero and failover begins. To avoid the countdown delay, highlight the monitored machine name in the Failover Control Center window and select **Failover**.

NOTE: Remove the source from the network before testing failover. If the source is still on the network, an IP address conflict will occur.

Failover manual intervention

The **Require User Intervention to Initiate Failover** option generates a prompt through the Failover Control Center. The prompt directs the administrator to click **OK** to proceed with the failover or **Cancel** to cancel the failover process. Failover can be initiated when appropriate from the Failover Control Center



NOTE: If the Failover Control Center is not running when a failure occurs, the manual intervention dialog box will appear the next time the Failover Control Center is started.

When a failure occurs, an alert is forwarded to the Windows Event Viewer. You can then start the Failover Control Center and respond to the Manual Intervention Required prompt.

If SNMP is installed and configured, an SNMP trap is also generated. When using a third-party SNMP manager, an e-mail or page can be generated to notify you of the failure.

Managing Failover Through the Text Client

Through the Failover Control Center, the failover process can be configured, monitored, and tested. In addition, the Failover Control Center can be used to initiate the failback process.

Configuring monitoring and failover

1. If you are using Storage Mirroring replication, establish a source/target connection by following the steps in [Establishing a new connection](#) on page 6-16.
2. Specify a target machine by using the `target` command.

Command	<code>TARGET</code>
Description	Identifies a machine as the active target machine
Syntax	<code>TARGET <target_machine></code>
Options	<code>target_machine</code> —Name of the machine
Examples	<code>target jersey</code>
Notes	You must be logged into a machine using the <code>login</code> command before using the <code>target</code> command.

3. Determine what NICs are available on the specified target by using the `niclist` command.

Command	<code>NICLIST</code>
Description	Displays the NICs available on the specified target machine. Each NIC is assigned an integer value and this value is used in the <code>monitor move</code> command.
Syntax	<code>NICLIST [target_machine]</code>
Options	<code>target_machine</code> —Name of the machine
Examples	<code>niclist jersey</code>
Notes	<ul style="list-style-type: none">◆ If you do not specify a machine name, the value from the current target will be returned. If you have not identified a target, no data will be returned.◆ If you have not logged into the target machine, no data will be displayed.

4. Establish a monitor to use with the remaining monitor commands by using the `monitor create` command.

Command	<code>MONITOR CREATE</code>
Description	Establishes a source machine as a failover monitor. This is the machine that will be monitored by a target machine in case it should experience a failure.
Syntax	<code>MONITOR CREATE <source_machine></code>
Options	<code>source_machine</code> —Name of the machine
Examples	<code>monitor create indy</code>

-
5. Specify that you want to use the monitor that was just created by using the `monitor use` command.

Command	<code>MONITOR USE</code>
Description	Specifies the source machine designated as the monitor that will be used in subsequent monitor commands
Syntax	<code>MONITOR USE <monitor></code>
Options	<code>monitor</code> —Name of the source machine designated as the monitor
Examples	<code>monitor use indy</code>

6. Configure the monitor settings (including the IP address to monitor, the target NIC that will assume the IP address when it fails, the monitor interval and missed packets) by using the `monitor move` command.

Command	<code>MONITOR MOVE</code>
Description	Designates the IP address that will be failed over to the specified target NIC.
Syntax	<pre>MONITOR MOVE <IP_address> TO NIC <target_NIC> USE NONE ICMP INTERVAL <interval> TIMEOUT <timeout> SERVICE INTERVAL <interval> TIMEOUT <timeout> [monitor] <NOTEST> [monitor]</pre>
Options	<ul style="list-style-type: none">◆ <code>IP_address</code>—The IP address which should be moved during failover◆ <code>target_NIC</code>—The integer value of the target NIC obtained from the <code>niclist</code> command◆ <code>NONE</code>—Does not actively monitor the source machine. You will be responsible for identifying when a failure has occurred and initiating failover manually.◆ <code>ICMP</code>—Monitors the source by using ICMP pings◆ <code>INTERVAL interval</code>—The frequency, in seconds, of the monitor requests sent to the source machine to see if it is online and active◆ <code>TIMEOUT timeout</code>—The number of seconds before failover will occur. This number is reset to its maximum each time the source sends a response to the monitor request.◆ <code>SERVICE</code>—Monitors the source by using UDP heartbeats.◆ <code>monitor</code>—Name of the source machine designated as the monitor◆ <code>NOTEST</code>—Allows you to failover an IP address without sending monitor requests or expecting responses from the source. This option should only be used if you are monitoring multiple IP addresses but do not want to send monitor requests to each address.
Examples	<ul style="list-style-type: none">◆ <code>monitor move 205.31.2.57 to nic 1 use icmp interval 5 timeout 25</code>◆ <code>monitor move 205.31.2.68 to nic 2 use none notest</code>
Notes	<ul style="list-style-type: none">◆ If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.◆ If you choose to use UDP heartbeats (the service option), your source and target must be using the same Storage Mirroring port settings. If you choose to use ICMP pings (the ICMP option), the port settings can be different.

```

Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)
Target Machine      Transmit      Mirror Replication
jersey              started        0      0
  1 marketing                idle replicating

=====
Command:

=====
> target jersey
> niclist jersey
> monitor move 206.31.4.57 to nic 1 use icmp interval 5 timeout 5

```

7. To remove an IP address from an established monitor, use the **monitor remove** command.

Command	MONITOR REMOVE
Description	Removes an IP address that is currently being monitored
Syntax	MONITOR REMOVE <IP_address> [<i>monitor</i>]
Options	<ul style="list-style-type: none"> ◆ <i>IP_address</i>—The currently monitored IP address that should be removed ◆ <i>monitor</i>—Name of the source machine designated as the monitor
Examples	<ul style="list-style-type: none"> ◆ <code>monitor remove 205.31.2.57 indy</code> ◆ <code>monitor remove 205.31.2.68</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

8. Configure the failover settings by using the `monitor option` command.

Command	MONITOR OPTION
Description	Configures the settings to determine how failover will be performed
Syntax	<code>MONITOR OPTION [MOVENAME NOMOVENAME] [MOVEADDRESSES NOMOVEADDRESSES] [MOVESHARES NOMOVESHARES] [ADD REPLACE] [, USESHAREFILE NOUSESHAREFILE][, FAILONE FAILALL] [, FODELAY NOFODELAY] [, FBDELAY NOFBDELAY] [, INTERVENTION NOINTERVENTION] [, FOHOSTNAME NOFOHOSTNAME] [, FBHOSTNAME NOFBHOSTNAME][<i>monitor</i>]</code>
Options	<ul style="list-style-type: none">◆ MOVENAME—Moves the server name during failover◆ NOMOVENAME—Does not move the server name during failover◆ MOVEADDRESSES—Moves the IP address(es) during failover◆ NOMOVEADDRESSES—Does not move the IP address(es) during failover◆ MOVESHARES—Moves the shares during failover◆ NOMOVESHARES—Does not move the shares during failover◆ ADD—Specifies the source machine's identity is added to the target machine's identity when failover occurs◆ REPLACE—Specifies that the source machine's identity replaces the target machine's identity when failover occurs◆ USESHAREFILE—Use the Storage Mirroring generated .shr file to determine shares◆ NOUSESHAREFILE—Do not use the Storage Mirroring generated .shr file to determine shares◆ FAILONE—When multiple IP addresses exist on a monitor machine, only the failed address is failed over to the target machine◆ FAILALL—When multiple IP addresses exist on a monitor machine, all of the addresses will fail over to the target machine even if only one address fails◆ FODELAY—Guarantees that the pre-failover script has completed before failing over◆ NOFODELAY—Does not guarantee that the pre-failover script has completed before failing over◆ FBDELAY—Guarantees that the pre-failback script has completed before failing back◆ NOFBDELAY—Does not guarantee that the pre-failback script has completed before failing back◆ INTERVENTION—Specifies that network administrator intervention is required before failover begins◆ NOINTERVENTION—Specifies that network administrator intervention is not required before failover begins◆ FOHOSTNAME—Removes the host SPN from Active Directory on the source and adds it to Active Directory on the target during failover◆ NOFOHOSTNAME—Does not modify Active Directory on the source or target during failover◆ FBHOSTNAME—Returns the host SPN on the source and target back to their original settings during failback◆ NOFBHOSTNAME—Does not modify Active Directory on the source or target during failback◆ <i>monitor</i>—Name of the source machine designated as the monitor
Examples	<ul style="list-style-type: none">◆ <code>monitor option</code>◆ <code>monitor option nomovename</code>
Notes	The default settings are <code>movename</code> , <code>moveaddress</code> , <code>moveshares</code> , <code>add</code> , <code>usesharefile</code> , <code>failall</code> , <code>fodelay</code> , <code>fbdelay</code> , <code>intervention</code> , <code>nofohostname</code> , and <code>nofbhostname</code> .


```

Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)
Target Machine      Transmit      Mirror Replication
jersey              started        0      0
  1 marketing              idle replicating

=====
Command:

=====
> niclist jersey
> monitor move 206.31.4.57 to nic 1 use icmp interval 5 timeout 5
> monitor option add, usesharefile, failall

```

9. If you specified to failover or failback the hostname in the `monitor option` command. You will need to specify a user and the associated password that has update privileges within Active Directory using the `monitor account add` command. This allows SPNs to be created and deleted during failover and failback.

Command `MONITOR ACCOUNT ADD`

Description Identifies a user account with update permissions within Active Directory

Syntax `MONITOR ACCOUNT ADD <username> <password> [target_machine]`

Options

- ◆ *username*—Name of the user in the format domain\user or user@domain
- ◆ *password*—Password associated with username
- ◆ *target_machine*—Name of the target machine

Examples `monitor account add administrator **** jersey`

Notes

- ◆ The monitor account add command is not available when scrolling through the Text Client command history.
- ◆ If characters in the password include non-alphanumeric characters, the password field must be enclosed in quotation marks.
- ◆ The password cannot be a Storage Mirroring keyword. These are any DTCL commands (source, target, and so on) or any DTCL shortcut commands (env, mon, rep, and so on).
- ◆ The password cannot be blank.

10. Specify any scripts by using the `monitor script add` command.

Command	<code>MONITOR SCRIPT ADD</code>
Description	Specifies the scripts that should be run during the failover and failback processes
Syntax	<code><u>MONITOR</u> SCRIPT ADD <type> <script_name> [<i>monitor</i>]</code>
Options	<ul style="list-style-type: none">◆ type—Any of the following script types:<ul style="list-style-type: none">◆ PREFAILOVER—Specifies that the file is a pre-failover script to be run on the target before failover occurs◆ POSTFAILOVER—Specifies that the file is a post-failover script to be run on the target after failover occurs◆ PREFAILBACK—Specifies that the file is a pre-failback script to be run on the target before failback occurs◆ POSTFAILBACK—Specifies that the file is a post-failback script to be run on the target after failback occurs◆ SRCPOSTFAILBACK—Specifies that the file is a post-failback script to be run on the source after failback occurs◆ script_name—Full path and name of the script file◆ monitor—Name of the source machine designated as the monitor
Examples	<code>monitor script add prefailback "C:\Program Files\OpenView\Storage Mirroring\preback.bat"</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

11. Remove any scripts by using the `monitor script remove` command.

Command	<code>MONITOR SCRIPT REMOVE</code>
Description	Specifies the scripts that should not be run during the failover and failback processes
Syntax	<code><u>MONITOR</u> <u>SCRIPT</u> <u>REMOVE</u> <type> [<i>monitor</i>]</code>
Options	<ul style="list-style-type: none">♦ type—Any of the following script types:<ul style="list-style-type: none">♦ PREFAILOVER—Specifies that the file is a pre-failover script to be run on the target before failover occurs♦ POSTFAILOVER—Specifies that the file is a post-failover script to be run on the target after failover occurs♦ PREFAILBACK—Specifies that the file is a pre-failback script to be run on the target before failback occurs♦ POSTFAILBACK—Specifies that the file is a post-failback script to be run on the target after failback occurs♦ SRCPOSTFAILBACK—Specifies that the file is a post-failback script to be run on the source after failback occurs♦ monitor—Name of the source machine designated as the monitor
Examples	<ul style="list-style-type: none">♦ <code>monitor script remove prefailback</code>♦ <code>monitor script remove postfailover</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

12. Review your failover settings by using the `monitor display` command.

Command	<code>MONITOR DISPLAY</code>
Description	Displays the monitoring and failover configuration settings for the specified monitor machine
Syntax	<code><u>MONITOR</u> <u>DISPLAY</u> <<i>monitor</i>></code>
Options	monitor —Name of the source machine designated as the monitor
Examples	<code>monitor display indy</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

```

Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)

- Monitor settings for indy -
Flags: ADD, USESHAREFILE, FAILALL
Target Pre-Failover script:c:\scripts\preover.txt
Target Post-Failover script:c:\scripts\postover.txt
Target Pre-Failback script:c:\scripts\preback.txt
Target Post-Failback script:c:\scripts\postback.txt
Source Post-Failback script:c:\scripts\src_postback.txt
    (press any key)

> monitor script add prefailback c:\scripts\preback.txt
> monitor script add postfailback c:\scripts\postback.txt
> monitor script add srcpostfailback c:\scripts\src_postback.txt
> monitor display

```

13. Start monitoring by using the `monitor start` command.

Command	MONITOR START
Description	The configured failover target starts monitoring the source machine that was designated as the monitor.
Syntax	<u>MONITOR</u> START <monitor> [ON <target_machine>]
Options	<ul style="list-style-type: none"> ♦ monitor—Name of the source machine designated as the monitor ♦ target_machine—Name of the target machine that will be monitoring the source machine
Examples	<code>monitor start indy on jersey</code>

14. If after starting a monitor you need to stop it, use the `monitor stop` command.

Command	MONITOR STOP
Description	Stops monitoring the source machine for a failure
Syntax	<u>MONITOR</u> STOP <monitor> [ON <target_machine>]
Options	<ul style="list-style-type: none"> ♦ monitor—Name of the source machine designated as the monitor ♦ target_machine—Name of the target machine that will no longer be monitoring the source machine
Examples	<code>monitor stop indy on jersey</code>

-
15. If you need to clear the monitor settings so that they can be reconfigured, use the `monitor clear` command.

Command	<code>MONITOR CLEAR</code>
Description	Clears all of the failover configuration and monitoring parameters for the specified monitor machine
Syntax	<code><u>MONITOR</u> CLEAR [<i>monitor</i>]</code>
Options	<i>monitor</i> —Name of the source machine designated as the monitor
Examples	<code>monitor clear indy</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

16. If you need to delete an established monitor, use the `monitor delete` command.

Command	<code>MONITOR DELETE</code>
Description	Deletes the specified failover monitor and all of its parameters
Syntax	<code><u>MONITOR</u> <u>DELETE</u> <<i>monitor</i>></code>
Options	<i>monitor</i> —Name of the source machine designated as the monitor
Examples	<code>monitor delete indy</code>
Notes	In order to successfully delete a monitor, the monitor must not be running on the server. Use the <code>monitor stop</code> command to ensure the monitor is not running.

Testing failover

Verify that the source machine is offline and not connected to the network to avoid IP address conflicts and trigger failover using the `failover` command.

Command	<code>FAILOVER</code>
Description	Manually initiates the failover process for the specified monitor machine
Syntax	<code>FAILOVER <monitor> [ON <target_machine>]</code>
Options	<ul style="list-style-type: none">♦ monitor—Name of the source machine designated as the monitor♦ target_machine—Name of the target machine
Examples	<code>failover indy on jersey</code>

Initiating failback

When failover occurs, a source machine has failed. The steps below must be completed in order to complete failback without IP address or name conflicts.

1. Verify the source machine is not connected to the network. If it is, disconnect it.
2. Resolve the source machine problem that caused the failure.

WARNING: Do not connect the source machine to the network at this time.

3. In the Text Client, login to the target machine that is currently standing in for the failed source by using the `login` command.
4. Identify the machine as the target by using the `target` command.

Command	TARGET
Description	Identifies a machine as the active target machine
Syntax	TARGET <target_machine>
Options	target_machine —Name of the machine
Examples	<code>target jersey</code>
Notes	You must be logged into a machine using the <code>login</code> command before using the <code>target</code> command.

5. Initiate failback by using the `failback` command. If you specified a pre-failback script in your failover configuration, that script will be executed at this time.

Command	FAILBACK
Description	Initiates the failback process for the specified monitor machine
Syntax	FAILBACK <monitor> [ON <target_machine>] [REMONITOR NOREMONITOR]
Options	<ul style="list-style-type: none">◆ monitor—Name of the source machine designated as the monitor◆ target_machine—Name of the target machine◆ REMONITOR—Automatically continues monitoring the source machine after failback◆ NOREMONITOR—Automatically discontinues monitoring the source machine after failback.
Examples	<code>failback indy on jersey</code>
Notes	If you do not select <code>remonitor</code> or <code>noremonitor</code> , you will be prompted after failback is complete to select whether or not to continue monitoring.

When failback is complete, you will be prompted to determine if you want to continue monitoring the source server. Bring the source back online and then enter `Y` or `N` to indicate if you want to continue monitoring the source machine. After you have selected whether or not to continue monitoring the source machine, the source post-failback script, if configured, will be started.

NOTE: At this time, the source machine must be online and Storage Mirroring must be running to ensure that the source post-failback script can be started. If the source has not completed its boot process, the command to start the script may be lost and the script will not be initiated.

Failback is now complete. If you need to restore newer files from the target to the source, see [Restoring Storage Mirroring Data Through the Text Client](#) on page I2-7.

12 ► Restoration


The Storage Mirroring restoration process provides an easy method for copying replicated data from the target back to its original location on the source. The process only requires you to select the source, target, and the appropriate replication set. There is no need to select files or to remember where the data came from on the source since that information is maintained by Storage Mirroring.

Restoration can be used if the source data is lost due to a disk crash or when the most up-to-date data exists on the target due to failover. At the time of a source machine failure, your Storage Mirroring target will contain the same data as your Storage Mirroring source. If you are using the Storage Mirroring failover capabilities, users can continue updating data on the target machine while the problems on the source are resolved. Because of the continued updates on the target, when the source machine is ready to come back online, the two machines will no longer contain the same data. Restoration is the process of copying the up-to-date data from the target back to the original source location when bringing the source back online.

Keep in mind the following points when completing a restoration:

- ◆ Do not let users access the source or target machines during the restoration process. The restoration process is similar to the original connection process, except that only mirroring is activated. Therefore, the data must not be changing on the target during the restore. Changes made to files on the target during the restore will not be transmitted back to the source. Additionally, files that are in use on the source may prevent the restore from updating those files.
- ◆ During a restoration, only the data and replication set definition are restored back to the source; shares are not created on the source. Shares that were created on the target during failover will need to be created on the source. In the case of a complete restore of the source, all shares will have to be re-created manually.
- ◆ If you are using Replicate NT Security by Name, before restoring from the target back to the source, you must enable that option on the target machine. Since a target machine acts like a source during a restoration, this option is enabled through the Server Properties Source tab.

How the Restoration Process Works



1. The network administrator resolves the source machine problems and in the process reinstalls Storage Mirroring, if necessary.
2. The network administrator initiates restoration through one of the Storage Mirroring clients.
3. Storage Mirroring determines which files to restore by using the replication set database backup from the target machine or the original replication set database on the source machine. This option is defined by the network administrator.

NOTE: If using the backup database on the target, you have the option of redirecting the restored data to a new source machine or to a different location on the original source machine.

4. Storage Mirroring determines which files from the replication set will be mirrored to the source machine by the user-specified conditions set up by the network administrator.
5. Storage Mirroring mirrors the data from the target to the source.
6. The source machine returns an acknowledgment to the target when each mirror packet is received.

Restoration Options

When restoring Storage Mirroring data, you must specify the source and target machines and the replication set from the source machine that contains the data that was lost and needs to be restored, as well as configure how the restoration process will occur.

- ◆ **Original Source**—If you are restoring to a Storage Mirroring server other than where a replication set was created, you must specify the original source machine.
- ◆ **Restore From**—You must specify the Storage Mirroring target machine where the copy of the data now resides and will be restored from.
- ◆ **Replication Set**—You must specify the replication set that was originally defined to send the data from the source to the target.
- ◆ **Restore To**—You must specify the Storage Mirroring machine where you want to restore the data. (In the Text Client, this option is called *original*.)
- ◆ **Use Backup Replication Set**—You have the option of using the backup replication set database stored on the target for the restoration or, if available, you can attempt to use the replication set from the source machine.

NOTE: If Use Backup Replication Set is selected during a restoration, Storage Mirroring attempts to use the backup of the replication set database located on the target machine to identify the files to be restored. The backup files are:

- ◆ `source_machine_name.db`
- ◆ `source_machine_name.xfp`

These files are located in the directory where Storage Mirroring is installed. If inconsistencies in the network name-resolution exist, the name of the `.DB` and `.XFP` files on the target may not be the name Storage Mirroring expects. When these files are not found, the restore fails. Prior to a restore, verify that the `.DB` and `.XFP` files on the target machine match the source machine name.

- ◆ **Restore Replication Set**—You have the option of restoring the target's copy of the replication set to the source machine during the restore process. This allows you to establish another connection after the restoration process without having to re-create the replication set.
- ◆ **Restore Conditionals**—You have three restore conditionals that can be applied during the restoration process:
 - ◆ **Overwrite existing files during restore**—With this option, the restoration process overwrites all existing files on the source machine and restores those files that do not exist.
 - ◆ **Only if backup copy is more recent**—With this option, the restoration process can be limited to only overwriting the files that are newer on the target than on the source machine. Without this option, all files are overwritten regardless of the file's date and time stamps.
 - ◆ **Use block checksum comparison**—With this option, the restoration process can be limited to only overwriting blocks of data that are different on the source than on the target as determined by the block checksum comparisons. With this option, only the blocks that are different are overwritten.

NOTE: During a restoration, temporary connections are established between the source and target machines. Other connections cannot be established between these two machines until the temporary connection is terminated. This may take up to two minutes after a restoration is complete. Any connections attempted prior to the temporary connection being terminated will fail. Any connections attempted after the temporary connection is terminated, will function properly.

When a file is restored, the global unique ID (GUID) used in Windows 200x is also restored from the target's attribute database.

The following table illustrates how the restoration conditionals work together, as well as how they work with the global checksum setting on the Source tab of the Server Properties. (Even though the target is doing the work, the settings on the Server Properties Source tab do influence how the target transmits the data back to the source.) Determine which options you need to select when configuring the restoration by identifying which type of action you want Storage Mirroring to perform during the restoration process.

Restoration Conditional Options

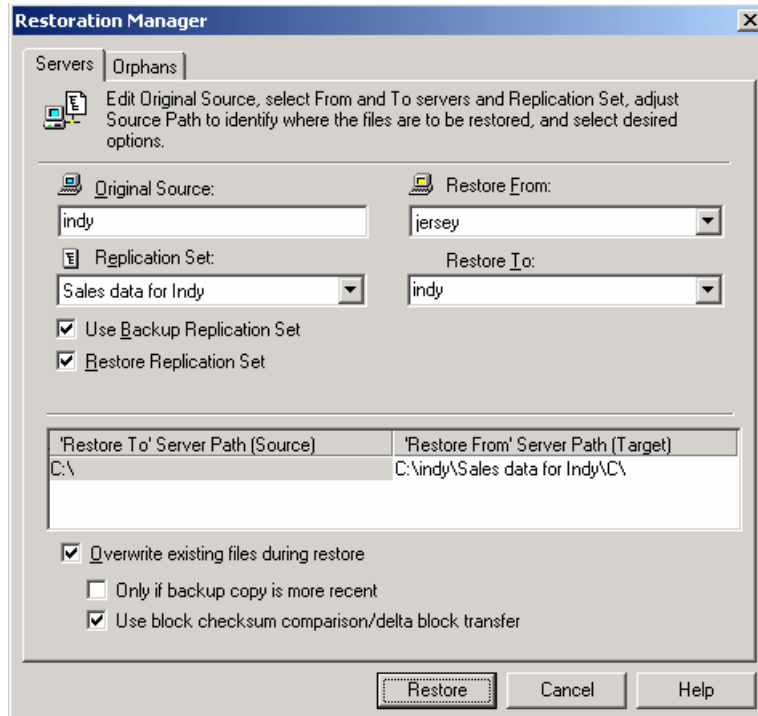
Block Checksum All Server Properties Source Tab ^a	Overwrite Files Restoration Option	Only if Backup is Newer Restoration Option	Block Checksum Restoration Option	Action Performed
(✓)				If no restoration options are selected, only the files that exist on the target but not on the source will be transmitted to the source.
(✓)	✓			All files on the target are restored to the source. The restoration sends the entire file.
(✓)	✓	✓		Any file that is newer on the target than on the source based on date and/or time is transmitted to the source. The restoration sends the entire file.
	✓		✓	Any file that is different on the target and source based on date, time, and/or size is flagged as different. The restoration then performs a checksum comparison on the flagged files and only sends those blocks that are different to the source.
✓	✓		✓	The restoration performs a checksum comparison on all files and only sends those blocks that are different to the source.
(✓)	✓	✓	✓	Any file that is newer on the target than on the source based on date and/or time is flagged as different. The restoration then performs a checksum comparison on the flagged files and only sends those blocks that are different to the source.

a. A checkmark enclosed in parentheses (✓) indicates that the global option on the Server Properties source tab can be on or off. The use of this option does not change the action performed during the restoration.

NOTE: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the Block Checksum All option on the Server Properties Source tab to ensure proper file comparisons.

Restoring Storage Mirroring Data Through the Management Console

1. From the Management Console, select **Tools, Restoration Manager**.



2. Select the **Original Source** machine. This is your source machine where the data originally resided.
3. Select the **Restore From** machine. This is the target machine where the backup data is stored.
4. **Replication Set** contains the replication set information stored on the target machine (the machine in **Restore From**). If no replication sets are available, the list will be blank. Select the replication set that corresponds to the data that you need to restore.

NOTE: Restoration applies to the entire replication set and should not be used for individual file restoration. Individual files can be restored manually by copying the files from the target to the source machine using native operating system tools.

5. Select the **Restore To** machine. This is the machine where the backup data will be sent.
6. The **Source Path** and **Target Path** will automatically be populated when the replication set is selected. The source path is the directory that is the common parent directory for all of the directories selected in the replication set. If the replication set crosses volumes, then there will be a separate source path for each volume. The target path is the path on the target server where the replicated files were sent.

NOTE: By default, the data will be returned to the same location it came from. However, Storage Mirroring will allow the data to be placed in a location other than the original.

Restoring across a NAT router requires the ports to be the same as the original connection. If the ports have been modified (manually or reinstalled), you must set the port numbers to the same values as the last valid source/target connection.

7. Select the **Use Backup Replication Set** check box to use the backup of the replication set stored on the target machine for the restore. If this check box is not marked, you will be accessing the replication set from the source machine.
8. Select the **Restore Replication Set** check box to restore the target's backup of the replication set database to the source machine during the restore process.

-
9. Select the restoration conditionals that you want to use. For details on how these conditionals work together, see the table [Restoration Conditional Options](#) on page 12-4.

NOTE: If you configure delete or move orphans, this process will be performed on the machine specified under **Restore To**.

The orphans configuration is not saved. If desired, you need to enable orphans each time you restore data.

10. If you want to configure orphan files, click the Orphans tab. For detailed information on configuring orphans, see [Configuring orphan files](#) on page 7-9.
11. Click **Restore** to begin the restoration. After the restoration is complete, the connection will automatically disconnect.

Restoring Storage Mirroring Data Through the Text Client

1. Log on to your source machine using the `login` command.

Command	<code>LOGIN</code>
Description	Log on to a Storage Mirroring machine
Syntax	<code>LOGIN <machine> <username> <password> [domain]</code>
Options	<ul style="list-style-type: none">◆ machine—Name of the machine◆ username—Name of the user◆ password—Password associated with username.◆ domain—If logging in using a domain account, this is the domain name. If logging in using a local account, this is the machine name.
Examples	<code>login indy administrator *****</code>
Notes	<ul style="list-style-type: none">◆ The <code>login</code> command is not available when scrolling through the Text Client command history.◆ If characters in the password include non-alphanumeric characters, the password field must be enclosed in quotation marks.◆ The password cannot be a Storage Mirroring keyword. These are any DTCL command (source, target, and so on.) or any DTCL shortcut command (env, mon, rep, and so on).

2. Log on to your target machine using the `login` command.
3. Identify your source by using the `source` command.

Command	<code>SOURCE</code>
Description	Identifies a machine as the active source machine
Syntax	<code><u>SOURCE</u> <source_machine></code>
Options	source_machine —Name of the machine
Examples	<code>source indy</code>

```
Source Machine: indy           Storage Mirroring version 4.4           (Access Level: FULL)

=====
Command: restore marketing from jersey ,overwritenewer, usetargetdb
=====
>restore marketing from jersey ,overwritenewer, usetargetdb
```

4. Restore Storage Mirroring data from target to the source by using the `restore from` command.

Command	<code>RESTORE</code>
Description	Initiates the restoration process
Syntax	<code>RESTORE <repset> FROM <target_machine> [ORIGINAL <original_source_machine>] [, OVERWRITE NOOVERWRITE][, OVERWRITENEWER NOOVERWRITENEWER] [, USETARGETDB NOSETARGETDB][, RESTOREDBTOO NORESTOREDBTOO] [, CHECKSUM NOCHECKSUM][, ORPHANS NOORPHANS]</code>
Options	<ul style="list-style-type: none">◆ <code>repset</code>—Name of the replication set◆ <code>target_machine</code>—Name of the target machine◆ <code>original_source_machine</code>—Specifies that the replication set was not originated on the current source machine◆ <code>OVERWRITE</code>—Overwrites files on the source◆ <code>NOOVERWRITE</code>—Does not overwrite files on the source◆ <code>OVERWRITENEWER</code>—Overwrites files on the source even if the source file is newer than on the target◆ <code>NOOVERWRITENEWER</code>—Does not overwrite files on the source that are newer on the source than on the target◆ <code>USETARGETDB</code>—Uses the replication set from the target machine◆ <code>NOSETARGETDB</code>—Uses the replication set from the source machine◆ <code>RESTOREDBTOO</code>—Restores the replication set from the target to the source◆ <code>NORESTOREDBTOO</code>—Does not restore the replication set from the target to the source◆ <code>CHECKSUM</code>—Performs a block checksum comparison and only restores those blocks that are different◆ <code>NOCHECKSUM</code>—Does not perform a block checksum comparison and restores those files that are different◆ <code>ORPHANS</code>—Removes orphan files on the target◆ <code>NOORPHANS</code>—Does not remove orphan files on the target
Examples	<code>restore "Exchange Repset" from jersey, overwritenewer, usetargetdb</code>
Notes	<ul style="list-style-type: none">◆ The default settings for this command are <code>overwrite</code>, <code>overwritenewer</code>, <code>usetargetdb</code>, <code>restoredbtoo</code>, and <code>noorphans</code>.◆ The source command is required before each use of the <code>restore</code> command.◆ If a path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.◆ If you are using the <code>restore</code> command with the DTCL interactive or batch client (<code>dtcl -i</code> or <code>dtcl -f</code>), the command requires the original source option. This option is not required if you are using the Text Client.

Automating Restore

You can use a post-failback script on the source to automate the restore process. For example, the following two files automate the restore process. In your failover configuration, identify `spostback.bat` as your post-failback script on the source machine. This batch file calls the second script below, `spostback.txt`, which automates the restoration.

SPOSTBACK.BAT

```
cd C:\Program Files\OpenView\Storage Mirroring
cmd /c DTCL -f "C:\Program Files\OpenView\Storage Mirroring\spostback.txt" >> "C:\Program
Files\OpenView\Storage Mirroring\restore.log"
```

NOTE: The second and third lines of the above batch file should appear on one line in your batch file. They appear on two lines in the text above so that you can see the entire command.

Spostback.bat Command	Description
cd C:\Program Files\OpenView\Storage Mirroring	Changes to the directory where the Storage Mirroring program files are located.
cmd /c DTCL -f "C:\Program Files\OpenView\Storage Mirroring\spostback.txt" >> "C:\Program Files\OpenView\Storage Mirroring\restore.log"	<p>Starts an instance of the command prompt, carries out the specified command, and then stops. In this case, the specified command is to run the Storage Mirroring Command Line client with the file name option using the <code>spostback.txt</code> file.</p> <p>The >> portion of the command outputs the results from the automated restore script file to the file specified, <code>restore.log</code>.</p>

SPOSTBACK.TXT

```
# SAMPLE SOURCE FAILBACK AND RESTORE SCRIPT                                     #
#                                                                              #
# THIS POST-FAILBACK SCRIPT EXECUTES ON THE SOURCE.                         #
# THIS SCRIPT PERFORMS A TARGET-SOURCE RESTORE AND                         #
# REESTABLISHES THE CONNECTION AND STARTS MIRRORING.                       #
# YOU NEED THE FOLLOWING INFORMATION TO EDIT THIS SCRIPT:                   #
#     target                                                                #
#     source                                                                #
#     repset                                                                #
#     username                                                             #
#     password                                                             #
#                                                                              #

# Login to source machine                                                    #
# Full access level will return 2, so exit on failure                      #
# Change <source><username><password> to match your server                  #
$level = LOGIN <source> <username> <password>;
if $level < 2 then
    WRITE "full access denied on <source>";
    WRITE $level;
    quit;
end

# Login to target machine                                                    #
# Full access level will return 2, so exit on failure                      #
# Change <target><username><password> to match your server                  #
$level = LOGIN <target> <username> <password>;
if level < 2 then
    WRITE "Full access denied on <target>";
    WRITE $level;
    quit;
end

# Successful source returns 0, so exit on failure                          #
$ret = SOURCE <source>;
if $ret < 0 then
    WRITE "Could not source <source>";
    WRITE $ret;
    quit;
end
```

SPOSTBACK.TXT (continued)

```
# Restore target data to source before                                #
# recreating original connection                                    #
# Restore ID should be 1 or greater, so exit on failure            #
$restID=RESTORE <repset> from <target> original <source> overwrite;
if $restID < 1 then
    WRITE "restore failed from <target>";
    WRITE $restID;
    quit;
end

SOURCE <source>;
WAITONRESTORE $restID;
# Add a hard wait of 20 seconds to clear queues.                    #
WAIT 20000;
# Successful source returns 0, so exit on failure                  #
$ret = SOURCE <source>;
if $ret < 0 then
    WRITE "Could not source <source>";
    WRITE $ret;
    quit;
end

# Re-create original connection                                    #
# Connection ID should be 1 or greater, so exit on failure        #
$connectionID = CONNECT <repset> to <target> MAP EXACT NOMIRROR;
if $connectionID < 1 then
    WRITE "Could not connect <repset> to <target>";
    WRITE $connectionID;
    quit;
end

# Successful mirror start returns 0, so exit on failure            #
$ret = MIRROR START $connectionID DIFFERENT, CHECKSUM;
if $ret < 0 then
    WRITE "Could not start mirror to <target>";
    WRITE $ret;
    quit;
end

WAITONMIRROR $connectionID;

WRITE "All operations completed successfully";
```

13 ► Monitoring Tools

Storage Mirroring offers a variety of tools for monitoring Storage Mirroring activity. With these monitoring tools, you can obtain the status of Storage Mirroring machines, connection statistics, and alerts. Alerts are processing notifications, warnings, and/or errors.

- ◆ **Management Console Tools**

- ◆ The Management Console main window displays source/target connection and statistical information in the right pane.
- ◆ At-a-glance monitoring includes icons and graphics to quickly determine the status of Storage Mirroring machines and connections.
- ◆ The Message Window displays alert messages.

- ◆ **Text Client Tools**

- ◆ The Text Client main window displays source/target connection and statistical information.
- ◆ The `status` command displays connection and status information.

- ◆ **Operating System Dependent Tools**

- ◆ The Windows Event Viewer receives alerts from the Storage Mirroring service and logs them to the application log.
- ◆ The Windows Performance Monitor displays statistical information including connection and processing statistics.

- ◆ **Operating System Independent Tools**

- ◆ DTStat is a utility that reports statistical information including connection and processing statistics.
- ◆ LogViewer is a utility that gathers Storage Mirroring alerts which are written to disk and can be displayed using the LogViewer utility.
- ◆ An SNMP Management Console receives both statistics and alerts from Storage Mirroring.

NOTE: Regardless of the monitoring tool you are using, you may see replication operations during a mirror, even if you are not currently replicating any data. These operations are pre-allocating space in order to avoid fragmentation on the target. These operations are normal.

Management Console Tools

The Management Console tools include:

- ◆ The Management Console main window which displays source/target connection and statistical information in the right pane.
- ◆ At-a-glance monitoring includes icons and graphics that indicate the status of Storage Mirroring machines and connections.
- ◆ The Message Window displays alert messages.

Management Console statistics

By highlighting a source on the left pane of the Management Console main window, the current connections and their statistics are displayed on the right. Scroll to the right to view the additional statistics. The following table lists each statistic available for the Management Console main window and what that statistic identifies.

The statistics marked with an asterisk (*) are not displayed by default. To change the display of the statistics, see [Customizing the statistics](#) on page 13-4.

Statistic	Description
Replication Set	Replication set indicates the name of the connected replication set.
Connection ID	The connection ID is the incremental counter used to number each connection established. This number is reset to one each time the Storage Mirroring service is restarted.
Target Name	The name of the target as it appears in the server tree in the left pane of the Management Console. If the server's name is not in the server tree, the IP address will be displayed.
Target IP	The target IP is the IP address on the target machine where the mirroring and replication data is being transmitted.
Target Status	<ul style="list-style-type: none">◆ OK—The target machine is active, online, and keeping up with data sent from the source machine.◆ Busy—The target machine is active and online, but is not keeping up with data that is sent from the source machine. The source machine is paused until the target catches up.◆ Not Loaded—The Storage Mirroring service is no longer running on the target machine or a TDU connection has been established.◆ Pause Pending—The target machine has been paused by user intervention. The pause is pending while the target completes the execution of operations that were already in progress.◆ Paused—The target machine has been paused by user intervention.◆ Restore Required—The target machine is or has been standing in for the source machine due to failover. So that data on the target is not overwritten by data on the source, which may be out-of-date, you must perform a restore back to the same source or override the setting by manually starting mirroring, replication, delete oprhans, or verification with a remirror.◆ Retrying—The target machine is retrying operations for that connection.◆ Error—The target machine has dropped operations because they could not be applied successfully after retrying. (This state will only be displayed if all components (source, target, and Management Console) are running version 4.4.1 or later.) <p>This field may not be updated until there is source/target activity.</p>
*Commit Mode	<p>The commit mode status indicates the connection status.</p> <ul style="list-style-type: none">◆ Real-time—Data is being transmitted to the target machine in real-time.◆ Scheduled—Data is waiting to be transmitted to the target machine until one or more transmit options have been met.

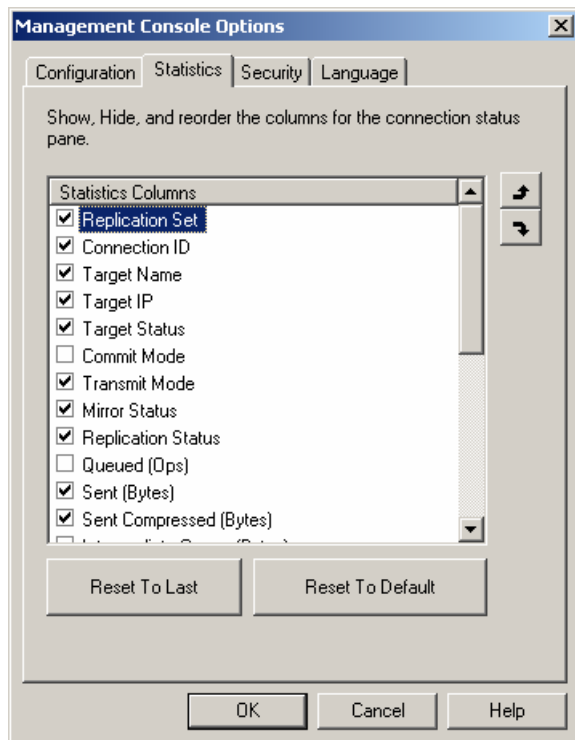
Statistic	Description
Transmit Mode	<ul style="list-style-type: none"> ◆ Started—Data is being transferred to the target machine. ◆ Paused—If the transmission is real-time and the transmission has been paused, the Transmit Mode indicates Paused. ◆ Scheduled—If the transmission is scheduled, the Transmit Mode indicates Scheduled. ◆ Stopped—Data is not being transferred to the target machine. ◆ Error—There is a transmission error.
Mirror Status	<ul style="list-style-type: none"> ◆ Mirroring—If the file size of the replication set has not been calculated and the data is being mirrored to the target machine, the Mirror Status will indicate Mirroring. ◆ Idle—Data is not being mirrored to the target machine. ◆ Paused—Mirroring has been paused. ◆ Percentage Complete—If the file size of the replication set has been calculated and the data is being mirrored to the target machine, the Mirror Status will display the percentage of the replication set that has been sent. ◆ Waiting—Mirroring is complete, but data is still being written to the target. ◆ Restoring—Data is being restored from the target to the source. ◆ Verifying—Data is being verified. ◆ Removing Orphans—Storage Mirroring is checking for orphan files within the target path location (files that exist on the target but not on the source). These files will be removed.
Replication Status	<ul style="list-style-type: none"> ◆ Replicating—Data is being replicated to the target machine. ◆ Ready—There is no data to replicate to the target machine. ◆ Stopped—Replication has stopped. ◆ Pending—If auto-remirror is enabled and you have experienced a source or target failure and recovery, the status will change to pending while the connections are reestablished and will update when the remirror begins. If auto-remirror is disabled and you have experienced a source or target failure and recovery, replication will be Pending until a remirror is performed. Without a remirror, data integrity cannot be guaranteed. ◆ Out of Memory—Kernel memory has been exhausted.
*Queued (Ops)	The queued (ops) statistic indicates the total number of mirror and replication operations that are in the source queue.
Sent (Bytes)	The sent (bytes) statistic indicates the total number of mirror and replication bytes that have been transmitted to the target.
Sent Compressed (Bytes)	The sent compressed (bytes) statistic indicates the total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as sent (bytes).
*Intermediate Queue (Bytes)	The intermediate queue (bytes) indicates the total amount of memory being used by the operations buffer queue.
Disk Queue (Bytes)	The disk queue (bytes) indicates the amount of disk being used to queue data on the source.
Queued Replication (Bytes)	The queued replication (bytes) statistic is the total number of replication bytes that are remaining to be transmitted from the source.
Sent Replication (Bytes)	The sent replication (bytes) statistic is the total number of replication bytes that have been transmitted to the target.
*Sent Compressed Replication (Bytes)	The sent compressed replication (bytes) statistic is the total number of compressed replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as sent replication (bytes).
*Queued Mirror (Ops)	The queue mirror (ops) statistic is the total number of mirror operations in the queue.
Sent Mirror (Bytes)	The sent mirror (bytes) statistic is the total number of mirror bytes that have been transmitted to the target.

Statistic	Description
*Sent Compressed Mirror (Bytes)	The sent compressed mirror (bytes) statistic is the total number of compressed mirror bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as sent mirror (bytes).
Skipped Mirror (Bytes)	The skipped mirror (bytes) statistic is the total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is not different on the source and target machines.
Remaining Mirror (Bytes)	The remaining mirror (bytes) statistic is the total number of mirror bytes that are remaining to be sent to the target.
* Queued Replication (Ops)	The queued replication (ops) statistic is the total number of replication operations in the queue.
Last File Touched	The last file touched identifies the last file that Storage Mirroring transmitted to the target.
Connected Since	Connected since is the date and time indicating when the current connection was made. This field is blank, indicating that a TCP/IP socket is not present, when the connection is waiting on transmit options or if the transmission has been stopped. This field will maintain the date and time, indicating that a TCP/IP socket is present, when transmission has been paused.

Customizing the statistics

You can change the statistics that are displayed on the Management Console main window by following these steps:

1. Select **File, Options** and click the **Statistics** tab. The statistics will be listed with check boxes to the left of each item.



2. Mark the check box to the left of each statistic that you want to appear on the Management Console main window.
3. Clear the check box to the left of each statistic that you do not want to appear on the Management Console main window.

Changing the order of the statistics

The statistics appear on the Management Console main window in the order they appear on the **Statistics** tab. There are two ways to change the order in which statistics are displayed:

- ◆ **Using the Mouse**—Highlight the statistic to be moved on the **Statistics** tab and select the up or down arrow button, to the right of the vertical scroll bar, to move the selection up or down in the list. Repeat this process for each statistic that needs to be moved until you reach the desired order.
- ◆ **Using the Keyboard**—Use the arrow keys to highlight a statistic on the **Statistics** tab and use the key combinations **ALT+UP ARROW** and **ALT+DOWN ARROW** to move the highlighted selection up and down, respectively. Repeat this process for each statistic that needs to be moved until you reach the desired order.

Saving statistics modifications

Click **OK** to automatically apply and save any changes that have been made to the order or display of the Management Console statistics.

Reverting to the last configuration

If you have made changes to the statistics list and have not yet saved them, you can go back to the previously used settings by clicking **Reset to Last**. This will revert the list back to the last saved settings.

Resetting the default configuration












To return the statistics list to the Storage Mirroring default selection and order, click **Reset to Default**. To view the default selection, see the chart under [Management Console statistics](#) on page 13-2.

At-a-Glance monitoring

At-a-glance monitoring includes icons and colors that indicate the status of Storage Mirroring machines and connections. Each icon or color that indicates the status of Storage Mirroring activity is described in the following tables.




- ◆ The icons in the left pane indicate what modules are loaded on each machine.

Left Pane Icons

Icon	Description
	A blue server in front of a yellow folder indicates a group.
	The group icon with a yellow triangle and a black exclamation point indicates there is a communication error with one of the servers in the group.
	The group icon with a red circle with a white X indicates there is a connection error with one of servers in the group.
	A blue machine indicates a Storage Mirroring source.
	A yellow machine indicates a Storage Mirroring target.
	Yellow and blue machines indicate a Storage Mirroring source and target.
	A single machine with a hammer indicates that neither the source nor target modules are loaded.
	A machine with two red vertical lines over it indicates that the target module is paused.
	A machine with red horizontal lines over it indicates that a restore may be required because the target machine is or has been standing in for the source machine due to failover.
	Any of the machine icons can appear with a red X. The red X indicates: <ul style="list-style-type: none">◆ The Management Console cannot communicate with that Storage Mirroring machine.◆ There is a problem with an established Storage Mirroring connection. Look in the right pane of the Management Console or log on to the machine in the left pane to determine the exact problem.
	A machine icon can appear with a black X. The black X indicates that the machine is not running Storage Mirroring.

- ◆ The icon to the left of the machine name indicates the security access that is granted for that machine. For detailed information on security and access rights, see [Security](#) on page 15-1.

Security Icons

Icon	Description	Access Granted
	This icon is a computer with a gear and it indicates the Storage Mirroring security is set to administrator access.	Administrator rights
	This icon is a computer with a magnifying glass and it indicates the Storage Mirroring security is set to monitor only access.	Monitor rights
	This icon is a lock and it indicates the Storage Mirroring security is set to no access.	No rights

- ◆ The right pane displays the following information depending on the item selected in the left pane of the Management Console and the tab selected on the right pane.







Right Pane Display




Left Pane	Right Pane	Description
Group	No Tabs	All groups and servers contained in the selected group are displayed.
Source Machine	Source Tab	The active connections, if any, from that source machine are displayed.
Source Machine	Target Tab	The display is blank.
Target Machine	Source Tab	The display is blank.
Target Machine	Target Tab	The active connections, if any, to that target machine are displayed. ^a The server name is also displayed, if it is in the Management Console tree. If it is not in the tree, the IP address is displayed.
Replication Set	No Tabs	The available target machines are displayed so that you can drag and drop a replication set onto a target to establish a connection.
Volume or Directory	No Tabs	The directories and files contained in that volume or directory are displayed.

a.If transmission is stopped (manually stopped, outside of a scheduled transmission window, network error between the source and target, and so on) the replication set will not be displayed on the Target tab. When transmission (re)starts, the replication set will (re)appear.

- ◆ When you have a group or server selected in the left pane of the Management Console, the icons displayed in the following table may be displayed in the right pane.



Right Pane Icons Associated with Groups and Servers

Icon	Description
	A blue server in front of a yellow folder indicates a group.
	The group icon with a yellow triangle and a black exclamation point indicates there is a communication error with one of the servers in the group. Drill down through the group until you find the server that has the communication error.
	The group icon with a red circle with a white X indicates there is a connection error with one of servers in the group. Drill down through the group until you find the server that has the connection error.
	There are no established Storage Mirroring connections to this machine. This icon also indicates that communications between the Management Console and the machine are working properly.
	Any established Storage Mirroring connections are working properly, as well as communications between the Management Console and the machine.
	<p>A problem exists with an established Storage Mirroring connection. The problem could be:</p> <ul style="list-style-type: none"> ◆ The transmission has been broken ◆ Replication is in a pending state ◆ The target is not loaded ◆ The target is busy ◆ The kernel memory reaches the Stop Replicating Limit <p>To determine the exact problem, log onto the source machine in the left pane of the Management Console and, using the horizontal scroll bar at the bottom of the right pane of the Management Console, scroll to find the statistic that appears in red. This statistic indicates the connection problem.</p>

Icon	Description
	<p>A problem exists with the communication between the Management Console and the Storage Mirroring machine. The problem could be:</p> <ul style="list-style-type: none"> ◆ The Storage Mirroring service has stopped ◆ The machine has crashed ◆ There are network problems ◆ The machine running the Management Console is experiencing network problems
	The machine is not running Storage Mirroring.
	A restore may be required because the target machine is or has been standing in for the source machine due to failover.

- ◆ The following table identifies the different icons and colors that you may see in the right pane of the Management Console when you have a machine with an established connection selected. These icons let you know, at-a-glance, the status of your Storage Mirroring connections.


Connection Icons

Icons and Colors	Description
	A green checkmark appears in the right pane of the Management Console to the left of the connection data line for an established connection that is working properly.
	<p>A red X appears in the right pane of the Management Console to the left of a connection that contains an error. There are five cases that can lead to an error state:</p> <ul style="list-style-type: none"> ◆ The transmission has been broken ◆ Replication is in a pending state ◆ The target is not loaded ◆ The target is busy ◆ The source machine has reached the Stop Replicating Limit indicating that it is out of memory <p>To determine the exact problem, use the horizontal scroll bar at the bottom of the right pane of the Management Console and scroll to find the connection data item that appears in red. The connection problem is indicated in red.</p> <p>Note: The red X will only appear if replication is in a pending state due to processing. If replication is manually stopped, either by not initiating it when the connection is established or by manually stopping it, the connection will have a green check mark even though replication is not occurring.</p>
White	If the background is white, the Management Console has received a response from the source indicating that the source and Management Console are communicating.
Gray	If the background is gray, the number of missed responses from the source has been exceeded, therefore, communications between the Management Console and the source has been lost. The connection statistics are not updating when the background is gray. The background will return to white and the statistics will begin updating again when the Management Console receives a response from the source machine. For information on configuring the missed responses, see Customizing when the background color and icons change on page 13-9.

Customizing when the background color and icons change

Communications between the Management Console and the source are monitored by requests and responses. The Management Console sends a request to the source and expects a response back. If the Management Console exceeds the number of responses that can be missed, the Management Console changes icons and color to alert you to potential problems. In the event that you expect some missed responses because of heavy network traffic, a slow network, busy server, and so on, you can configure when the Management Console icons and colors change.

1. From the Management Console, select **File, Options**.



On the Configuration tab, you will see **Site Monitor** and **Connection Monitor** with the similar choices. The **Site Monitor** is monitoring your Storage Mirroring site - those icons in the world view and server tree. The **Connection Monitor** is monitoring each specific connection - the statistics and icons that are displayed in the right pane when a machine is highlighted in the left pane. These two separate monitoring capabilities allow for flexible monitoring of each individual connection, separate from the monitoring of the site as a whole.

The **Check Status Interval** allows you to specify, for the entire site, how often the Management Console sends out requests to the source(s). The **Missed Status Responses**, which can be configured separately for the site and for connections, is the number of responses from the source that can be missed before the icons and color will change.

NOTE: If the **Site Monitor** and **Connection Monitor** settings are different, at times, the icons and color may not be synchronized between the left and right panes.

2. Under **Site Monitor**, specify **Check Status Interval** to identify the number of seconds between requests sent from the Management Console to the source in order to update the site at-a-glance monitoring. Valid values are between 0 and 3600. The default setting is 30 seconds.
3. Under **Site Monitor**, specify **Missed Status Responses** to identify the number of responses from the source that can be missed before the Management Console considers communications lost and updates the icons in the site at-a-glance monitoring. Valid values are between 1 and 100. The default setting is 2.
4. Under **Connection Monitor**, specify **Missed Status Responses** to identify the number of responses from the source that can be missed before the Management Console considers communications lost and updates the connection monitoring pane. Valid values are between 0 and 1000. The default setting is 5.
5. Click **OK** to save the settings.

Management Console message window

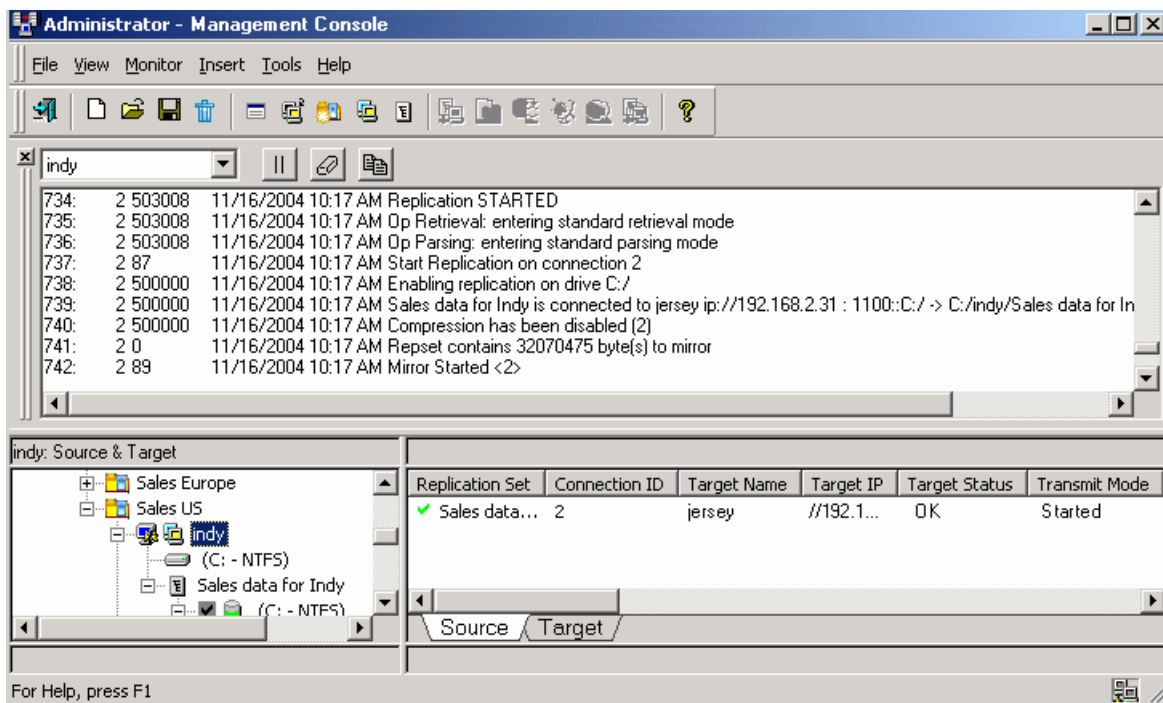
The Management Console message window displays Storage Mirroring processing alerts. These alerts include notifications, warnings, and errors. Use the instructions below to create and control the Management Console message window.

NOTE: The standard appearance of the message window is a white background. If your message window has a gray background, the window is inactive. The Management Console may have lost communications with that server, for example, or you may no longer be logged into that server.

The message window is limited to the most recent 1000 lines. If any data is missing (the message window was paused while activity occurred, the Storage Mirroring service was restarted, and so on) an entry in red will indicate the missing data. Regardless of the state of the message window, all data is maintained in the Storage Mirroring log on the server.





I. Open a new message window by using any of the following methods:

- ◆ Right-click on the machine that you want to monitor in the left pane and select **New, Message Window**.
- ◆ Select the Message Window icon from the toolbar.
- ◆ Select **Monitor, New Message Window** and identify the **Server** that you want to monitor.



NOTE: Repeat this step to open multiple message windows.

-
2. To control the window after it is created, use either of the following methods to access the control methods listed in the table:
 - ◆ Right-click on the message window and select the appropriate control.
 - ◆ Select the appropriate toolbar control
 - ◆ Select **Monitor**, the name of the message window, and the appropriate control.

Message Window Control	Description	Toolbar Icon
Close	Closes the message window	
Clear	Clears the message window	
Pause/ Resume	<p>Pauses and resumes the message window.</p> <ul style="list-style-type: none">◆ Pausing prevents new messages from being displayed in the message window so that you are not returned to the bottom of the message window every time a new message arrives. The messages that occur while the window is logged are still logged to the Storage Mirroring log file. See LogViewer on page 13-39 for additional details, including log file limits.◆ Resuming displays the messages that were held while the window was paused and continues to display any new messages. <p>Pausing is automatically initiated if you scroll up in the message window. The display of new log messages will automatically resume when you scroll back to the bottom.</p>	
Copy	Allows you to copy selected text	
Options	This control is only available from the Monitor menu. Currently, there are no filter options available so this option only allows you to select a different server. In the future, this control will allow you to filter which messages to display.	

3. To change which server you are viewing messages for, select a different machine from the dropdown list on the toolbar. If necessary, the login process will be initiated.
4. To move the message window to other locations on your desktop, click and drag it to another area or double-click it to automatically undock it from the Management Console.

Text Client Tools

The text client tools include:

- ◆ The Text Client main window which displays source/target connection and statistical information
- ◆ DTCL status command which displays connection and status information

Text Client main window

After a source/target connection is established, the Text Client main window displays basic connection and statistical information for the active source machine.

NOTE: Use the `source` command to set a machine as the active source.

The Command Line Client does not display connection and statistical information.

Storage Mirroring version 4.4				
Source Machine: indy		(Access Level: FULL)		
Target Machine		Transmit	Mirror	Replication
jersey		started	15345	2568
1 marketing			idle	replicating
=====				
Command:				
=====				

Column	Description
Target Machine	<ul style="list-style-type: none">◆ The first row of the first column identifies the name of the target machine. You may also see (Paused) or (Restore Required) following the target machine name. This indicates that the target has been manually paused or a restore is required because of a failover. You must resume the target manually to clear the paused indicator. To clear the restore indicator, you must perform a restore back to the same source or override the setting by manually starting mirroring, replication, delete oprhans, or verification with a remirror.◆ The second row of the first column identifies the connection ID and the replication set name.
Transmit	The status of the transmission
Mirror	<ul style="list-style-type: none">◆ The first row of the third column displays the number of bytes that have been mirrored.◆ The second row of the third column identifies the status of the mirror process.
Replication	<p>The first row of the fourth column displays the number of bytes that have been replicated.</p> <p>The second row of the fourth column identifies the status of the replication process.</p>

DTCL status command

Use the `status` command to display statistical and informational data.

Command	<code>STATUS</code>
Description	Requests connection and statistical information
Syntax	<code>STATUS <u>C</u>ONNECT <u>M</u>IRROR <u>R</u>EPLICATE <u>T</u>RANSMIT <con_id></code>
Options	<ul style="list-style-type: none">◆ <code>CONNECT</code>—Displays connection information◆ <code>MIRROR</code>—Displays mirroring information◆ <code>REPLICATE</code>—Displays replication information◆ <code>TRANSMIT</code>—Displays transmission information◆ <code>con_id</code>—Connection ID assigned to the source/target connection
Examples	<ul style="list-style-type: none">◆ <code>status connect 1</code>◆ <code>status replicate 1</code>

```
Storage Mirroring version 4.4
Source Machine: indy (Access Level: FULL)

- Status for connection 1 -
Time established:          09:59:28 01/15/2005
Replication set:          Marking
Mirror bytes trans/queue: 12354      175
Replication bytes trans/queue: 6825      985
    (press any key)

> status connect 1
```

Operating System Dependent Tools

The operating system dependent tools include:

- ◆ Windows Event Viewer
- ◆ Windows Performance Monitor

Windows Event Viewer

In the Windows operating system, an event is a significant occurrence in the system or in an application that requires administrators to be notified. The operating system writes notifications for these events to a log file that can be displayed using the Windows Event Viewer. Three different log files are generated: system, security, and application logs. Storage Mirroring related events are sent to the application log and are available for viewing through the Event Viewer.

1. To access the Event Viewer, select **Programs, Administrative Tools (Common), Event Viewer**.
2. To view the events logged for applications, select **Log, Application**. The following information is displayed for each event in the Event Viewer window:
 - ◆ **Date**—The date the event occurred.
 - ◆ **Time**—The time the event occurred.
 - ◆ **Source**—The software that logged the event, which can be either an application or a component of the system, such as a driver.
 - ◆ **Category**—A classification of the event.
 - ◆ **Event**—Shows an event number to identify the specific event. The Event helps product-support representatives track events in the system.
 - ◆ **User**—Specific text that exactly matches text in the User name field. This field is not case sensitive.
 - ◆ **Computer**—The exact name of the computer where the logged event occurred. This field is not case sensitive.
3. To view additional details, double-click an event. The additional information is displayed in the Event Detail screen:
 - ◆ **Type**—A classification of the event by the Windows operating system, such as Error, Warning, Information, Success Audit, or Failure Audit.
 - ◆ **Description**—A description of the event.

NOTE: For additional information on customizing the Event Viewer (such as sorting the display, filtering the display, and so on), see your Windows reference guide or the Windows online help.

Event descriptions

The following table identifies the events generated by Storage Mirroring.

Category	ID	Severity	Message	Required Response
Activation Key	1	Error	This evaluation period has expired. Mirroring and replication have been stopped. To obtain a license, please contact your vendor.	Contact your vendor to purchase either a single or site license.
	2	Info	The evaluation period expires in %1 day(s).	Contact your vendor before the evaluation period expires to purchase either a single or site license.
	3	Info	The evaluation period has been activated and expires in %1 day(s).	Contact your vendor before the evaluation period expires to purchase either a single or site license.
	4	Warning	Duplicate activation codes detected on machine %1 from machine %2.	If you have an evaluation license or a site license, no action is necessary. If you have a single license, you must purchase either another single license or a site license.
	5	Error	This product edition can only be run on Windows 2000 Server or Advanced Server running the Server Appliance Kit.	Verify your activation code has been entered correctly and contact technical support.
	4100	Error	Product activation code is invalid. Please check that it is typed correctly and is valid for the version of the operating system in use.	If you are in the process of installing Storage Mirroring, verify that you are using a 16 character alpha-numeric code. If Storage Mirroring is already installed, confirm that the code entered is correct. If the code appears to be correct, contact technical support.
	4101	Error	This service will not run on this device. Contact your sales representative for upgrade procedures.	The activation code does not match the type of server you are attempting to run on. Contact your vendor for a new activation code or contact technical support.
ARP	5400	Info	Broadcasted new MAC address %1 for IP address %2.	No action required.
DbiTake	12288	Error	Incurred an exception accessing shared replication buffer. The last Word in the Data Window is the exception code.	Reboot the server and contact technical support if this event occurs again.
	12289	Error	Incurred an exception accessing one of the application's buffers. The last Word in the Data Window is the exception code.	Reboot the server and contact technical support if this event occurs again.
	16384	Error	Encountered an unrecoverable error. Replication is stopped. A reboot of the server is required to re-enable replication.	Reboot the server and contact technical support if this event occurs again.
	16385	Info	Internal function ReferenceOrBuild StreamInfo returned an unexpected result. The last Word in the Data Window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
	16386	Info	A very large write was logged. The last Word in the Data Window is the number of bytes written. Please save this event log and contact technical support.	Contact technical support.

Category	ID	Severity	Message	Required Response
DbiTake	16387	Info	A FILE_OBJECT collision was detected and managed. This message is for informational purposes only and does not indicate a problem.	No action required.
	16388	Info	Unable to buffer for name of replication pagefile. Replication is not affected. The paging of replication ops is disabled.	Reboot the server and contact technical support if this event occurs again.
	16389	Info	Unable to create a pagefile for replication. Replication is not affected. The paging of replication ops is disabled. The last Word in the Data Window is the NT status code.	The Storage Mirroring driver could not create a pagefile. Verify that the QJournalDir location is an NTFS volume.
	16390	Info	Unable to allocate pagein buffer for paging of replication. Replication is not affected. The paging of replication ops is disabled.	An allocation for the buffer used to read pagefile operations failed. Reboot the server and contact technical support if this event occurs again.
	16391	Error	An error occurred while accessing the replication pagefile. Replication has been stopped. The paging of replication ops is disabled. The last Word in the Data Window is the NT status code.	This error may be caused because the drive is out of space, is not an NTFS partition, is possibly going bad, and so on. Contact technical support.
	16392	Info	An internal error occurred. The last Word in the Data Window indicates the internal error code.	Contact technical support if this event occurs again.
	16393	Info	Failed to build data structure for "%2" because of insufficient OS stack space. This file will not be replicated.	If the specified file is in the replication set, reboot the server and contact technical support if this event occurs again. If the file is not contained in the replication set, this message can be disregarded.
	16394	Error	Dblhook detected a memory error, possibly caused by faulty hardware or a bad driver. The last Word in the Data Window indicates the internal error code.	This message may be caused by another application attempting to access memory reserved by Storage Mirroring. Contact technical support if this event occurs again.
	16395	Warning	Dblhook detected a hard link for a file. The volume GUID and file ID are stored in the Data Window. This file will not be replicated.	Files with hard links are not supported.
DTCounters	16396	Info	The driver made more than one pass before finding a valid buffer. The last Word in the Data Window indicates the number the tries.	Contact technical support if this event occurs again.
	1000	Error	An exception occurred: %l	Run the Storage Mirroring installation and select Repair. Contact technical support if this event occurs again.
	1001	Error	The Double-Take counter DLL could not initialize the statistics handler object to gather performance data.	Run the Storage Mirroring installation and select Repair. Contact technical support if this event occurs again.
	1002	Error	The Double-Take counter DLL could not map shared memory file containing the performance data.	Run the Storage Mirroring installation and select Repair. Contact technical support if this event occurs again.
	1003	Error	The Double-Take counter DLL could not open the "Performance" key in the Double-Take section of the registry.	Run the Storage Mirroring installation and select Repair. Contact technical support if this event occurs again.

Category	ID	Severity	Message	Required Response
DTCounters	1004	Error	The Double-Take counter DLL could not read the "First Counter" value under the Double-Take\Performance Key.	Run the Storage Mirroring installation and select Repair. Contact technical support if this event occurs again.
	1005	Error	The Double-Take counter DLL read the "First Help" value under the Double-Take\Performance Key.	Run the Storage Mirroring installation and select Repair. Contact technical support if this event occurs again.
	1006	Error	The Double-Take counter DLL could not create event handler for the worker thread.	Run the Storage Mirroring installation and select Repair. Contact technical support if this event occurs again.
Failback	5200	Info	Failback completed for %1.	No action required.
	5201	Info	IP address %1 was removed from target machine's %2 adapter.	No action required.
	5202	Error	Unable to Failback properly because IP address %1 was missing a corresponding SubNet Mask.	Contact technical support.
Failover	5100	Info	Failover completed for %1.	No action required.
	5101	Info	IP address %1 with subnet mask %2 was added to target machine's %3 adapter.	No action required.
	5102	Warning	%1 has reached a failover condition. A response from the user is required before failover can take place.	User intervention has been configured. Open the Failover Control Center and accept or decline the failover prompt.
	5103	Info	Started adding drive shares from %1 to %2.	No action required.
	5104	Info	%1 drive shares were taken over by %2.	No action required.
	5105	Info	Attempting to run the %1 script.	No action required.
	5106	Info.	The %1 script ran successfully.	No action required.
	5107	Error	Error occurred in running %1 script.	Verify that the script identified exists with the proper permissions.
	5108	Error	The source machine %1 is not responding to a ping.	This occurs when all monitored IP addresses on the source machine stop responding to pings. Countdown to failover will begin at the first occurrence and will continue until the source machine responds or until failover occurs.
	5109	Error	The public NIC on source machine %1 is not responding to a ping.	The failover target did not receive an answer to its ping of the source machine. Eventually, a failover will result. Investigate possible errors (down server, network error, etc.).
Monitoring	5300	Info	The following IP address was added to target's monitoring list: %1	No action required.
	5301	Info	The following IP address was removed from target's monitoring list: %1	No action required.
	5302	Info	Drive share information for %1 has been updated on the target machine.	No action required.
NSIPHDriver	9000	Error	Could not allocate memory for storing registry path	Reboot the server and contact technical support if this event occurs again.
	9001	Error	Thread creation failed	Reboot the server and contact technical support if this event occurs again.

Category	ID	Severity	Message	Required Response
NSIPHDriver	9002	Error	Could not allocate memory for parameter key	Reboot the server and contact technical support if this event occurs again.
	9003	Error	Could not read PlaceholderCount from the NSIPH Parameters key	Reboot the server and contact technical support if this event occurs again.
	9004	Error	Could not allocate buffer to hold initial IP address list for adapter	Reboot the server and contact technical support if this event occurs again.
	9005	Error	Could not allocate buffer to hold initial subnet mask list for adapter	Reboot the server and contact technical support if this event occurs again.
	9006	Error	Could not read registry value for IP address list or subnet mask list for adapter	Reboot the server and contact technical support if this event occurs again.
	9007	Error	Could not write registry value for IP address list for adapter	Reboot the server and contact technical support if this event occurs again.
	9008	Error	Could not write registry value for subnet mask list for adapter	Reboot the server and contact technical support if this event occurs again.
	9009	Error	Could not form full registry path to adapter	Reboot the server and contact technical support if this event occurs again.
Recognizer	12288	Error	Nsi-recognizer was started improperly. Nsi-recognizer must be started during boot. Change Nsi-recognizer's Startup value to Boot in the Devices applet of Control Panel, then reboot the system.	This error is no longer used.
	12289	Error	An internal error has occurred. Please report this error to the technical support department.	Contact technical support.
Resources	8192	Warning	Kernel memory was exhausted. Replication is stopped.	Contact technical support if this event occurs again.
	8192	Error	Failed to create device list entry due to insufficient resources.	Contact technical support if this event occurs again.
	8193	Error	Failed to create Transfer Thread. The last Word in the Data Window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
	8194	Error	Failed to create filename work queue. The last Word in the Data Window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
	8195	Error	Failed to create file close queue. The last Word in the Data Window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
	8196	Warning	The maximum amount of memory for replication queuing has been reached. Replication is stopped and memory is being freed.	Contact technical support if this event occurs again.
	8197	Error	Failed to create file delete queue. The last Word in the Data Window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
	8198	Error	Failed to save the registry path. The last Word in the Data Window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
	8199	Error	Failed to create tracked file queue. The last Word in the Data Window is the NT status code.	Reboot the server and contact technical support if this event occurs again.

Category	ID	Severity	Message	Required Response
Resources	8200	Warning	Failed to allocate buffer for querying very long filenames, i.e., names over 260 characters. Replication will continue for standard length filenames. The last Word in the Data Window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
	8201	Warning	Reached stop replicating page limit. Replication is stopped.	Configure the applications on the server to use less memory or add more memory so that the server does not get into a low memory state.
	8202	Error	Failed to create Tracked File System worker thread. The last Word in the Data Window is the NT status code.	The Storage Mirroring driver cannot create the thread used by the tracked file system. If a thread cannot be created, the system is low on critical resources. Reboot the server and contact technical support if this event occurs again.
	8203	Error	Failed to allocate a copy of the logfile path. The last word in the Data Windows is the NT status code.	Reboot the server and contact technical support if this event occurs again.
Service	3000	Info.	Logger service was successfully started.	No action required.
	3001	Info.	Logger service was successfully stopped.	No action required.
	4000	Info	Kernel was successfully started.	No action required.
	4001	Info	Target service was successfully started.	No action required.
	4002	Info	Source service was successfully started.	No action required.
	4003	Info	Source service was successfully stopped.	No action required.
	4004	Info	Target service was successfully stopped.	No action required.
	4005	Info	Kernel was successfully stopped.	No action required.
	4006	Error	Service has aborted due to the following unrecoverable error: %l	Restart the Storage Mirroring service.
	4007	Warning	Auto-disconnecting from %l (%2) for Replication Set %3, ID: %4 due to %5	<p>The connection is auto-disconnecting for one of the following reasons:</p> <ol style="list-style-type: none"> 1. The disk-based queue on the source has been filled 2. The service has encountered an unknown file ID 3. The target server has restarted 4. An error has occurred during disk queuing on the source (for example, Storage Mirroring cannot read from or write to the transaction log file) 5. An error has occurred during disk queuing on the target (for example, Storage Mirroring cannot read from or write to the transaction log file) <p>Contact technical support if this event occurs frequently.</p>
	4008	Info.	Auto-disconnect has succeeded for %l (%2) for Replication Set %3, ID: %4	No action required.

Category	ID	Severity	Message	Required Response
Service	4009	Info	Auto-reconnecting Replication Set %1 to %2 (%3)	No action required.
	4010	Info.	Auto-reconnect has succeeded connecting Replication Set %1 to %2 (%3)	No action required.
	4011	Error	Auto-reconnect has failed connecting Replication Set %1 to %2 (%3)	Manually reestablish the replication set to target connection.
	4014	Info	Service has started network transmission.	No action required.
	4015	Info	Service has stopped network transmission.	No action required.
	4016	Info	Service has established a connection to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4017	Info	Service has disconnected from %1 (%2) for Replication Set %3, ID: %4	No action required.
	4018	Warning	%1, however, mirroring and replication have been disabled as a restore is required due to a previous failover.	Perform a restoration.
	4019	Info	Service has started a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4020	Info	Service has paused a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4021	Info	Service has resumed a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4022	Info	Service has stopped a mirror to %1 for Replication Set %2, ID: %3, %4	No action required.
	4023	Info.	Service has completed a mirror to %1 %2 for Replication Set %3, ID: %4, %5	No action required.
	4024	Info	Service has started Replication to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4025	Info	Service has stopped Replication to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4026	Info	The target has been paused due to user intervention.	No action required.
	4027	Info	The target has been resumed due to user intervention.	No action required.
	4028	Warning	Registration of service class with Active Directory failed. Verify that the Active Directory server is up and the service has the proper permissions to update its entries.	Verify that the Active Directory server is running and that the Storage Mirroring service has permission to update Active Directory.
	4029	Warning	Registration of service instance with Active Directory failed. Verify that the Active Directory server is up and the service has the proper permissions to update its entries.	Verify that the Active Directory server is running and that the Storage Mirroring service has permission to update Active Directory.
	4030	Error	RSResource.dll has an unknown error. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

Category	ID	Severity	Message	Required Response
Service	4031	Error	RSResource.dll could not be opened. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4032	Error	The RSResource.dll component version does not match the component version expected by the product. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4033	Error	RSResource.dll build version is invalid. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4034	Error	Error verifying the service name. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4035	Error	Error verifying the product name. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4036	Error	Error verifying the vendor name. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4037	Error	Error verifying the vendor URL name. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4038	Error	Error verifying the product code. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4039	Error	Error while reading RSResource.dll. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4040	Error	The product code is illegal for this computer hardware. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4041	Error	The product code is illegal for this operating system version. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4042	Error	The product code requires installing the Windows Server Appliance Kit. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
	4043	Error	This product can only be run on a limited number of processors and this server exceeds the limit. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

Category	ID	Severity	Message	Required Response
Service	4044	Error	An error was encountered and replication has been stopped. It is necessary to stop and restart the service to correct this error.	Contact technical support if this error persists.
	4045	Error	%1 value must be between 1025 and 65535. Using default of %2.	Verify that the Storage Mirroring port value you are trying to use is within the valid range. If it is not, it will automatically be reset to the default value.
	4046	Error	This service failed to start because of a possible port conflict. Win32 error: %1	Verify that the Storage Mirroring ports are not conflicting with ports used by other applications.
	4047	Error	Could not load DynaZip DLL %1. The service will not start.	Reinstall the software, using the installation Repair option, to install a new copy of the DynaZip.dll Contact technical support if this error persists.
	4048	Info.	Service has started a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4049	Info.	Service has paused a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4050	Info.	Service has resumed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4051	Info.	Service has stopped a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4052	Info.	Service has completed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4053	Info.	Service has started a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4054	Info.	Service has paused a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4055	Info.	Service has resumed a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4056	Info.	Service has stopped a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4057	Info.	Service has completed a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4058	Info.	Service has started a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4059	Info.	Service has paused a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4060	Info.	Service has resumed a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4061	Info.	Service has stopped a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.

Category	ID	Severity	Message	Required Response
Service	4062	Info.	Service has completed a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
	4110	Warning	Target cannot write %1 due to target disk being full. Operation will be retried (%2 times or forever)	The disk on the target is full. The operation will be retried according to the TGExecutionRetryLimit setting.
	4111	Warning	Target can not write %1 due to a sharing violation. Operation will be retried (%2 times or forever)	A sharing violation error is prohibiting Storage Mirroring from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
	4112	Warning	Target can not write %1 due to access denied. Operation will be retried (%2 times or forever)	An access denied error is prohibiting Storage Mirroring from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
	4113	Warning	Target can not write %1 due to an unknown reason. Operation will be retried (%2 times or forever). Please check the log files for further information on the error.	An unknown error is prohibiting Storage Mirroring from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
	4120	Info	Target write to %1 was completed successfully after %2 retries.	No action required.
	4150	Error	Target write %1 failed after %2 retries and will be discarded. See the event log or log files for error conditions. After correcting the problem, you should re-mirror or run a verify to resynchronize the changes.	The operation has been retried according to the TGExecutionRetryLimit setting but was not able to be written to the target and the operation was discarded. Correct the problem and remirror the files.
	4155	Warning	The service was unable to complete a file system operation in the allotted time. See the log files for error conditions. After correcting the problem, remirror or perform a verification with remirror to synchronize the changes.	Correct the file system error and then remirror or perform a verification with remirror to synchronize the changes.
	4200	Info	In band task %1 submitted from %2 by %3 at %4	No action required.
	4201	Warning	In band task %1 discarded (submitted from %2 by %3 at %4)	A task may be discarded in the following scenarios: all connections to a target are manually disconnected, replication is stopped for all connections to a target, or an auto-disconnect occurs. If one of these scenarios did not cause the task to be discarded, contact technical support.
	4202	Info	Running %1 in band script: %2 (task %3 submitted from %4 by %5 at %6)	No action required.
	4203	Info.	Completed run of in band script: %1 (exit code %2)	No action required.
	4204	Error	Error running in band script: %1	Review the task and its associated script(s) for syntax errors.
	4205	Warning	Timeout (%1 seconds) running in band script: %2	The timeout specified for the script to complete has expired. Storage Mirroring will continue with normal processing. You may need to manually terminate the script if it will never complete.

Category	ID	Severity	Message	Required Response
Service	4206	Warning	Run timeout disabled for in band script: %1	The timeout period was set to zero (0). Storage Mirroring will not wait for the script to complete before continuing. No action is required.
	4207	Warning	In band scripts disabled by server - no attempt will be made to run %1	Enable task command processing.
	4300	Error	A connection request was received on the target before the persistent target paths could be loaded.	You may need to disconnect and reconnect your replication set.
	4301	Error	Unable to block target paths, the driver is unavailable.	If you need to block your target paths, contact technical support.
	4302	Info.	Target Path %1 has been successfully blocked	No action required.
	4303	Warning	Blocking of target path: %1 failed. Error Code: %2	If you need to block your target paths, contact technical support.
	4304	Info.	Target Path %1 has been successfully unblocked	No action required.
	4305	Warning	Unblocking of target path: %1 failed. Error Code: %2	If you need to unblock your target paths, contact technical support.
	4306	Warning	Target paths for source %1 (%2) Connection id: %3 are already blocked	No action required.
	4307	Warning	Target paths for source %1 (%2) Connection id: %3 are already unblocked	No action required.
	4308	Error	Error loading target paths for blocking, registry key %1 has been corrupted.	If you need to block your target paths, contact technical support.
	5000	Info.	Server Monitor service was successfully started.	No action required.
	5001	Info.	Server Monitor service was successfully stopped.	No action required.
	5002	Info	Placeholders were modified to %1.	No action required.
	5500	Warning	Could not connect to e-mail server. Check to make sure the SMTP server %1 is available (error code: %2).	Storage Mirroring could not connect to your SMTP server or the username and/or password supplied is incorrect. Verify that SMTP server is available and that you have identified it correctly in your e-mail notification configuration. Also verify that your username and password have been entered correctly.
	5501	Warning	E-mail notification could not be enabled (error code: %1).	This alert occurs if there is an unexpected error enabling e-mail notification during service startup. Check to see if any other errors related to e-mail notification have been logged. Also, check to make sure the Windows Management Instrumentation (WMI) service is enabled. If neither of these apply, contact technical support.
	5502	Warning	E-mail notification could not be initialized. Check to make sure Internet Explorer 5.0 or later is installed.	E-mail notification requires Internet Explorer 5.0 or later. Verify that you have this version or later installed on the Storage Mirroring server.

Category	ID	Severity	Message	Required Response
Service	5503	Warning	E-mail notification could not be processed. Check to make sure the correct version of SMTPMail.DLL is registered on the system (error code: %1).	If you are using Windows NT 4.0, e-mail notification requires Windows Management Instrumentation (WMI) to be installed. Verify that you have it installed on the Storage Mirroring server.
	5504	Warning	Could not load LocalRS.dll (for e-mail notification).	This alert occurs if there is an error loading the resource DLL for the service. Typically, this is caused by a missing LocalRS.dll file. Reinstall the software, using the installation Repair option, to install a new copy of the LocalRS.dll. Contact technical support if this error persists.
	5505	Warning	E-mail could not be sent. Check e-mail settings (error code: %1).	Verify that the e-mail server that you have identified in your e-mail notification configuration is correct.
	5506	Warning	One or more required e-mail settings have not been specified (error code: %1).	At a minimum, you must specify the e-mail server, the From and To addresses, and at least one type of event to include.
	5507	Warning	E-mail notification could not be initialized. Check to make sure WMI is installed and available (error code: %1).	If you are using Windows NT 4.0, e-mail notification requires Windows Management Instrumentation (WMI) to be installed. Verify that you have it installed on the Storage Mirroring server.
	5508	Warning	An error occurred connecting to the WMI namespace. Check to make sure the Windows Management Instrumentation service is not disabled (error code %1).	This alert occurs if there is an error with the Windows Management Instrumentation (WMI) service. Verify that you have it installed on the Storage Mirroring server and that it is enabled.
	5600	Warning	Part or all of the e-mail setting %1 is not in a valid format.	Verify that the include categories and exclude ID list are identified and formatted correctly. Make changes through the Text Client or the Management Console.
	10000	Error	Connect failed to node %1 for resource %2. Adding node to reconnect list.	Ensure that GeoCluster is running on all possible owners and that it can communicate on the network selected for mirroring and replication traffic. GeoCluster will try to reestablish a connection using the check unresponsive node interval specified for the resource.
	10001	Info	Reconnect succeeded to node %1 for resource %2. Will be added as a possible owner when mirror is complete.	No action required.
	10002	Error	Disk check failed on node %1 for resource %2. Removing as a possible owner.	Ensure that GeoCluster is running on all possible owners and that it can communicate on the public network. Also ensure that the disk specified for the resource is functioning correctly on all possible owners.

Category	ID	Severity	Message	Required Response
Service	10003	Error	Owner %1 of the quorum resource %2 couldn't access the arbitration path %3. Network may be down.	Ensure that the network used to access the arbitration path is up and that the server is operational. Also ensure that the arbitration share path does exist and that the account running the cluster service has write privileges to the share path.
	10004	Warning	Failover of the group %1 is being delayed. Group will be brought online when the target queue is below the limit or the timeout has expired.	No action required.
	10005	Info	Node %1 is taking ownership of the group %2. The group will be brought online on this node.	No action required.
System	4096	Warning	Unknown registry parameter specified. Parameter is named "%2."	Delete the parameter and report this issue to technical support.
	4097	Error	Failed to load nsi-filter. The last Word in the Data Window is the NT status code.	Contact technical support.
	4097	Error	IoRegisterFsRegistrationChange() failed. The last Word in the Data Window is the NT status code.	Contact technical support.
	4098	Error	Failed to create device "%2." The last Word in the Data Window is the NT status code.	Contact technical support.
	4098	Error	Failed to create a device. The last Word in the Data Window is the NT status code.	Contact technical support.
	4099	Error	Failed to register with nsi-filter. The last Word in the Data Window is the NT status code.	Contact technical support.
	4099	Error	Failed to attach to a device. The last Word in the Data Window is the NT status code.	Contact technical support.
	4100	Error	Failed to open file. The Status and IoStatusBlock status are in the data window.	Contact technical support.
	4100	Warning	Failed to obtain a name for a recognizer device object. The second to the last Word in the Data Window is a pointer to the device object. The last Word in the Data Window is the NT status code.	Contact technical support.
	4101	Info	An error occurred while trying to identify the system volume. Failure to identify the system volume may result in a slight impact to replication performance. The error code is the last word in the data window.	Contact technical support.
	4101	Error	RcgMountCompletion was called at a raised IRQL. Please report this error to the technical support department. The last Word in the Data Windows is the IRQL.	Contact technical support.

Windows Performance Monitor

Performance Monitor is a graphical tool for measuring the performance of your own computer or other computers on the network. It provides charting, alerting, and reporting capabilities that reflect both current activity and ongoing logging. You can open, browse, and chart log files later as if they reflected current activity. Storage Mirroring statistics are available through the Performance Monitor.

1. To access the Performance Monitor, select **Programs, Administrative Tools, Performance**.
2. Add data to monitor by right-clicking and selecting **Add** or using the add button from the toolbar.
3. Choose one of the following Performance Objects: **Double-Take Connection**, **Double-Take Kernel**, **Double-Take Security**, **Double-Take Source**, or **Double-Take Target**.
4. Depending on the monitor you have chosen, you may or may not have to select specific statistics to display from the chosen object. Select the statistic you want to monitor and click **Add**.

NOTE: If you are using the Alert monitor, specify an **Alert If** threshold.

5. You can switch from working with current values for current activity (real-time data) to viewing and manipulating existing log files. See Performance Monitor help for details. The default setting is current activity.

NOTE: For additional information and details on the Performance Monitor, see your Windows reference guide or the Windows online help.

Performance Monitor statistics

Object	Statistic	Description
Connection	Bytes in disk queue	The number of bytes in the source disk queue
	Bytes in replication queue	The number of replication bytes in the source queue
	Bytes in mirror queue	The number of mirror bytes in the source queue
	Bytes received	The number of bytes received by the target since the last Performance Monitor refresh
	Bytes transferred	The number of bytes transmitted from the source
	Compressed bytes transferred	The number of compressed bytes transmitted from the source
	Operations in acknowledgement queue	The number of operations waiting in the source acknowledgement queue
	Operations in command queue	The number of operations waiting in the source command queue
	Operations in mirror queue	The number of mirror operations in the source queue
	Operations in replication queue	The number of replication operations in the source queue
	Operations received	The number of operations received by the target since the last Performance Monitor refresh
	Operations resent	The number of operations re-sent since the last time the Storage Mirroring service was restarted on the source
	Operations transmitted	The number of operations transmitted from the source
	Task commands queued	The number of task commands queued on the source
	Task commands submitted	The number of task commands submitted on the source
	Tasks failed	The number of task commands that have failed to execute on the source

Object	Statistic	Description
Connection	Tasks ignored	The number of task commands that have been ignored on the source
	Tasks succeeded	The number of task commands that have succeeded on the source
Kernel	Activation code failures	The number of activation code failures when loading the source or target, since the last time the Storage Mirroring service was restarted on the source
	Double-Take queue memory usage	The amount of system memory in use by the Storage Mirroring queue
	Failed mirror operations	The number of mirror operations on the source that failed due to an error reading the file from the disk
	Failed replication operations	The number of replication operations on the source that failed due to an error reading the file from the disk
	Mirror Kbytes generated	The number of mirror kilobytes transmitted from the source
	Mirror operations generated	The number of mirror operations transmitted from the source
	Open Target Handles	The number of handles currently open on the target.
	Replication Kbytes generated	The number of replication kilobytes generated on the source by the file system driver
	Replication operations generated	The number of replication operations generated on the source by the file system driver
Security	Failed logins	Number of failed login attempts since the last time the Storage Mirroring service was restarted
	Successful logins	Number of successful login attempts since the last time the Storage Mirroring service was restarted
Source	Auto disconnects	The number of automatic disconnects since the last time the Storage Mirroring service was restarted on the source
	Auto reconnects	The number of automatic reconnects since the last time the Storage Mirroring service was restarted on the source
Target ^a	Bytes in Disk Queue	The number of bytes in the target disk queue
	Bytes in Queue	The number of bytes in the system memory and disk queues
	Mirror operations received	The number of mirror operations received on the target
	Operations received	The number of operations received on the target
	Ops Dropped	The number of operations dropped on the target since the last time the Storage Mirroring service was restarted on the target
	Ops Remaining	The number of operations on the target remaining to be applied
	Orphan Bytes	The number of orphan bytes removed from the target
	Orphan Directories	The number of orphan directories removed from the target
	Orphan Files	The number of orphan files removed from the target

Object	Statistic	Description
Target ^a	Retries	The number of retries performed on the target since the last time the Storage Mirroring service was restarted on the target
	Tasks failed	The number of task commands that have failed on the target.
	Tasks ignored	The number of task commands that have been ignored on the target
	Tasks succeeded	The number of task commands that have succeeded on the target

a.If you have multiple IP addresses connected to one target server, you will see multiple Target statistic sections for each IP address.

Operating System Independent Tools

The operating system independent tools include:

- ◆ DTStat
- ◆ LogViewer
- ◆ SNMP

DTStat

Statistics logging is the process of taking snapshots of Storage Mirroring statistical data and writing it to a file for future use. Statistics are automatically gathered and written to disk by the Storage Mirroring service. The DTStat file can be configured from both the Management Console and text clients and any configuration changes are detected and applied immediately without restarting the Storage Mirroring service.

DTStat log file

The statistics log file created is a binary file and can be configured through the Management Console, text clients, or from the command prompt. To view the log file, you must run DTStat from the command prompt.

Sample DTStat Log File

```
=====
01/14/05 12:48:04:2340
=====
IQALLOCATOR::Total Bytes: 0
SECURITY::Logins : 3 FailedLogins : 0
KERNEL::SourceState: 2 TargetState: 1 Start Time: Fri Jan 11 15:30:17 2005
  RepOpsGenerated: 4 RepBytesGenerated: 0
  MirOpsGenerated: 33164 MirBytesGenerated: 1083527492
  FailedMirrorCount: 0 FailedRepCount: 0 ActFailCount: 0
TARGET::Thread Count: 1 Ops Received: 33166 Mirror Ops Received: 33131
  Retries: 0 OpsDropped: 0 Ops Remaining: 0
  Orphan Files Removed: 0 Orphan Directories Removed: 0 Orphan Bytes Removed: 0
  Bytes In Target Queue: 0
SOURCE::autoDisConnects : 0 autoReConnects : 0
  lastFileTouched : E:/sqldata
CONNECTION:: IP_Address: 169.254.10.1 connectTime: Fri Jan 11 15:56:02 2005
  conState: 1 conOpsInCmdQueue: 0 conOpsInAckQueue: 0
  conOpsInRepQueue: 0 conOpsInMirQueue: 0 conBytesInRepQueue: 0
  conOpsTx: 33168 conBytesInMirQueue: 0 conBytesTx: 1083527492
  conOpsRx: 33168 conBytesRx: 1083527492 conResentOpCount: 0
  conBytesSkipped: 0 conMirrorBytesRemain: 0 conMirrorPercent: 100.0%
```

Modifying the log file settings through the Management Console

1. Right-click a machine in the left pane of the Management Console and select **Properties**.
2. Select the **Logging** tab.

The screenshot shows the 'Server Properties' dialog box with the 'Logging' tab selected. The dialog has a title bar with a close button. Below the title bar are two rows of tabs: 'General', 'Setup', 'Network', 'Queue', 'Source' in the first row, and 'Target', 'Database', 'Logging', 'E-mail Notification' in the second row. The 'Logging' tab is active. It contains several sections: 'Folder:' with a text box showing 'C:\Program Files\OpenView\Storage Mirroring' and a browse button (...); 'Messages & Alerts' with 'Maximum Length (bytes):' set to '1048576' and 'Maximum Files:' set to '5'; 'Verification' with 'Filename:' set to 'DTVerify.log', 'Maximum Length (bytes):' set to '1048576', a checked 'Append' checkbox, and 'Language:' set to 'English'; and 'Statistics' with 'Filename:' set to 'statistic.sts', 'Maximum Length (bytes):' set to '10485760', and 'Write Interval (minutes):' set to '5'. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

3. At the top of the tab, specify the **Folder** where the log files for messages, alerts, verification, and statistics will be saved.
4. Under Statistics, specify the following information:
 - ◆ **Filename**—The name of the statistics log file. The default file name is `statistic.sts`.
 - ◆ **Maximum Length**—The maximum length of the statistics log file. The default maximum length is 10 MB. Once this maximum has been reached, Storage Mirroring begins overwriting the oldest data in the file.
 - ◆ **Write Interval**—The frequency in which Storage Mirroring writes the statistical data to the statistics log file. The default is every 5 minutes.
5. Click **OK** to save the settings.

Modifying the log file settings through the text clients

1. To determine the current location where the log file is stored, use the `get` command with the `StatsFileName` option.

Command	<code>GET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	<code>GET <option> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>get netport</code>♦ <code>get VerifyLogName</code>♦ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. To modify the statistics log file name, use the `set` command with the `StatsFileName` option. The default is `statistic.sts`, but you can use any valid file name.

Command	<code>SET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	<code>SET <option>=<value> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>set netport=1100</code>♦ <code>set VerifyLogName="HA server.log"</code>♦ <code>set DefaultAddress="129.12.34.3"</code>
Notes	<ul style="list-style-type: none">♦ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed.♦ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

3. To determine the maximum size of the statistics log file, use the `get` command with the `StatsMaxFileSize` option.
4. To modify the maximum size of the statistics log file, use the `set` command with the `StatsMaxFileSize` option. The default is 10485760 (10 MB).
5. To determine the current interval for writing to the statistics file, use the `get` command with the `StatsWriteInterval` option.
6. To modify the interval for writing to the statistics file, use the `set` command with the `StatsWriteInterval` option. The default is 5 minutes. Valid values are between 0 and 65535.

Running DTStat

You can run DTStat from the Storage Mirroring Text Client or from any command prompt.

Running DTStat from the text clients

To determine the status of DTStat or to stop or start it through the Text Client, use the `statslog status`, `statslog start`, or `statslog stop` commands.

◆ Starting DTStat

Command	<code>STATSLOG START</code>
Description	Starts the DTStats statistics logging utility.
Syntax	<code>STATSLOG START <machine> [TO <filename>] [EVERY <minute>] [MAXSIZE <kilobytes>]</code>
Options	<ul style="list-style-type: none">◆ <i>machine</i>—Name of the machine◆ <i>filename</i>—Any valid path and filename to which the statistical information will be logged◆ <i>minutes</i>—The number of minutes to write to the log file◆ <i>kilobytes</i>—The maximum file size in kilobytes
Examples	<code>statslog start indy to dtstat.sts every 20 maxsize 200000</code>
Notes	<ul style="list-style-type: none">◆ The default settings are <code>statistic.sts</code>.◆ The default interval is every 5 minutes.◆ The default file size is 10485760 bytes (10 MB).

◆ Stopping DTStat

Command	<code>STATSLOG STOP</code>
Description	Stops the DTStats statistics logging utility.
Syntax	<code>STATSLOG STOP <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>statslog stop indy</code>

◆ Determining the status of DTStat

Command	<code>STATSLOG STATUS</code>
Description	Identifies whether or not the DTStats statistics logging utility is currently running
Syntax	<code>STATSLOG STATUS <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>statslog status indy</code>

Running DTStat from the command prompt

From a command prompt, use the **DTStat** command from the location where Storage Mirroring is installed.

Command	DTSTAT
Description	Starts the DTStats statistics logging utility from a command prompt.
Syntax	DTSTAT [-p][-i <interval>][-t <filename>] [-f <filename>][-s <filename>] [-st <filename>][-IP <address>][-START <mm/dd/yyyy hh:mm>][-STOP <mm/dd/yyyy hh:mm>] [-SERVER <ip_address> <port_number>]
Options	<ul style="list-style-type: none">◆ -p—Do not print the output to the screen◆ -i interval—Refresh from shared memory every <i>interval</i> seconds◆ -t filename—Save the data from memory to the specified binary file <i>filename</i>◆ -f filename—Reads from a previously saved binary file, <i>filename</i>, that was generated using the -t option instead of reading from memory◆ -s filename—Saves only the connection data from the data in memory to an ASCII, comma-delimited file, <i>filename</i>◆ -st filename—Saves only the target data from the data in memory to an ASCII, comma-delimited file, <i>filename</i>◆ -f filename1 -s filename2—Saves only the connection data from a previously saved binary file, <i>filename1</i>, to an ASCII, comma-delimited file, <i>filename2</i>◆ -f filename1 -st filename2—Saves only the target data from a previously saved binary file, <i>filename1</i>, to an ASCII, comma-delimited file, <i>filename2</i>◆ -IP address—Filters out the specified <i>address</i> in the IP address field and prints only those entries. Specify more than one IP address by separating them by a comma.◆ -START mm/dd/yyyy hh:mm—Filters out any data prior to the specified date and time◆ -STOP mm/dd/yyyy hh:mm—Filters out any data after the specified date and time◆ SERVER ip_address port_number—Connects DTStat to the specified IP address using the specified port number instead of to the local machine
Examples	<ul style="list-style-type: none">◆ DTStat -i 30◆ DTStat -p -i 5 -t dtstat.log◆ DTStat -f dtstat.log -s dtstat.csv -start 02/02/2005 09:25◆ DTStat -server 206.31.4.51 1106
Notes	<ul style="list-style-type: none">◆ This command is not case-sensitive.◆ If no options are specified, DTStat will print the output to the screen at an interval of every second.◆ If the statistics are not changing, DTStat will discontinue writing until statistics begin updating again.

DTStat statistics

The following statistics appear in the `statistic.sts` file. If you convert that file to an ASCII, comma-delimited file using the `dtstat -s` option, the heading labels may appear slightly different.

Category	Statistic	Description
Date/Time Stamp		The date and time that the snapshot was taken. This is the date and time that each statistic was logged. By default, these are generated once a second, as long as there are statistics being generated. If mirroring/replication is idle, then DTStat will be idle as well.
System Allocator		The number of bytes currently allocated to the system pagefile
IQ Allocator		The number of bytes currently allocated to the intermediate queue
Security	Logins	The number of successful login attempts
	FailedLogins	The number of failed login attempts
Kernel	SourceState	0—Source is not running 1—Source is running without the replication driver 2—Source is running with the replication driver
	TargetState	0—Target is not running 1—Target is running
	Start Time	Date and time stamp indicating when the Storage Mirroring service was loaded
	RepOps Generated	The number of replication operations generated by the file system driver. An op is a file system operation. Storage Mirroring replicates data by sending the file system operations across the network to the target. RepOpsGenerated indicates the number of file system operations that have been generated by replication.
	RepBytes Generated	The number of replication bytes generated by the file system driver. This is the number of bytes generated during replication. In other words, this is roughly the amount of traffic being sent across the network that is generated by replication. It does not take into account TCP/IP overhead (headers and such).
	MirOps Generated	The number of mirror operations transmitted to the target. Mirroring is completed by transmitting the file system operations necessary to generate the files on the target. This statistic indicates the number of file system operations that were transmitted during the initial mirror. It will continue to increase until the mirror is complete. Any subsequent remirrors will reset this field to zero and increment from there.
	MirBytes Generated	The number of mirror bytes transmitted to the target. This is the number of bytes generated during mirroring. In other words, this is roughly the amount of traffic being sent across the network that is generated by the mirror. It does not take into account TCP/IP overhead (headers and such). Again, any subsequent remirror will reset this field to zero and increment from there.
	FailedMirror Count	The number of mirror operations that failed due to an error reading the file from the disk
	FailedRep Count	The number of replication operations that failed due to an error reading the file from the disk
	ActFailCount	The number of activation code failures when loading the source or target. Activation codes can be bad for reasons such as: expiration of evaluation codes, duplicate codes, incorrect codes, etc.
	TargetOpen Handles	The number of handles currently open on the target.
Target ^a	PeerAddress	The IP address of the source machine

Category	Statistic	Description
Target ^a	LocalAddress	The IP address of the target machine.
	Ops Received	The total number of operations received by this machine as a target since the Storage Mirroring service was loaded
	MirrorOps Received	The total number of mirror operations received by this machine as a target since the Storage Mirroring service was loaded. This number does not reset to zero for remirrors.
	Retries	The number of retries performed before all operations were completed
	OpsDropped	The number of operations skipped during a difference mirror. During a difference mirror, if Storage Mirroring detects that there have been no changes to a file, then it will indicate the number of operations it did not send for this file in this field.
	Ops Remaining	The total number of operations that are left in the target queue
	OrphanFiles Removed	The number of orphan files removed from the target machine
	Orphan Directories Removed	The number of orphan directories removed from the target machine
	OrphanBytes Removed	The number of orphan bytes removed from the target machine
	BytesInTarget Queue	The number of bytes currently in the system memory queue on the target
	BytesInTarget Disk Queue	The number of bytes currently in the disk queue on the target
	Tasks Succeeded	The number of task commands that have succeeded on the target
	TasksFailed	The number of task commands that have failed on the target
	TasksIgnored	The number of task commands that have been ignored on the target
Source	auto DisConnects	The number of automatic disconnects since starting Storage Mirroring. Auto-disconnects occur because the source no longer sees the target. This could be because the connection between the two has failed at some point or because the target machine data is changing on the source faster than the source can get the data to the target. This field tracks the number of times an auto-disconnect has occurred since the Storage Mirroring service was started.
	autoRe Connects	The number of automatic reconnects since starting Storage Mirroring. Auto-reconnect occurs after a target machine is back online. This field tracks the number of times an auto-reconnect has happened since the Storage Mirroring service was started.
	lastFileTouched	The last filename that had a replication operation executed
Connection	conPeer Address	The IP address of the target machine
	connectTime	The time that this connection was established

Category	Statistic	Description
Connection	conState	<p>The state of the active connection</p> <p>0—None. This indicates a connection has not been established. Statistics are still available for the source and target machines.</p> <p>1—Active. This indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.)</p> <p>2—Paused. This indicates a connection that has been paused.</p> <p>4—Scheduled. This indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations, time frame limitations, and so on).</p> <p>8—Error. This indicates a connection that is not transmitting because something has gone wrong (for example, lost connection).</p> <p>Only the Scheduled and Error states can coexist. All other states are mutually exclusive. DTStat will display a conState of 12 when the connection is in both a scheduled and an error state because this is the sum of the two values (4 + 8).</p>
	conOpsInCmd Queue	The number of operations waiting to be executed on the target
	conOpsInAck Queue	The number of operations waiting in the acknowledgement queue. Each operation that is generated receives an acknowledgement from the target after that operation has been received by the target. This statistic indicates the number of operations that have yet to receive acknowledgement of receipt.
	conOpsInRep Queue	The number of replication operations currently waiting to be executed on the target
	conOpsInMir Queue	The number of mirror operations currently waiting to be executed on the target
	conBytesInRep Queue	The number of replication bytes remaining to be transmitted to the target
	conOpsTx	The number of operations transmitted to the target. This is the total number of operations that Storage Mirroring has transmitted as a source. In other words, the cumulative number of operations transmitted by this source to all connected targets.
	conBytesInMir Queue	The number of mirror bytes remaining to be transmitted to the target
	conBytesTx	The number of bytes transmitted to the target. This is the total number of bytes that Storage Mirroring has transmitted as a source. In other words, the cumulative number of bytes transmitted by this source to all connected targets.
	conBytes CompressedTx	The number of compressed bytes transmitted to the target.
	conOpsRx	The number of operations received by the target. The number of operations that the target for this connection (as indicated by the IP address field) has received from this source.
	conBytesRx	The number of bytes received by the target. The number of bytes that the target for this connection (as indicated by the IP address field) has received from this source.
	conResentOp Count	The number of operations resent because they were not acknowledged
	conBytesInDisk Queue	The number of bytes in the source disk queue
	conBytes Skipped	The number of bytes skipped during a difference mirror. During a difference mirror, if Storage Mirroring detects that there have been no changes to a file, then it will indicate the number of bytes it did not send for this file in this field.
	conMirrorBytes Remain	The number of mirror bytes remaining to be transmitted

Category	Statistic	Description
Connection	conMirrorPercent	The percentage of the mirror that has been completed. This field is determined if the replication set size was calculated.
	conTaskCmdsSubmitted	The number of task commands that have been submitted on the source
	conTaskCmdsQueued	The number of task commands that have been queued on the source
	conTasksSucceeded	The number of task commands that have succeeded on the source
	conTasksFailed	The number of task commands that have failed on the source
	conTasksIgnored	The number of task commands that have been ignored on the source

a.If you have multiple IP addresses connected to one target server, you will see multiple Target statistic sections for each IP address.

LogViewer

LogViewer is a utility that gathers Storage Mirroring alerts. These alerts, which are processing notifications, warnings, and error messages, are written to disk and can be displayed using the LogViewer.

Sample LogViewer Output

01/15/2005	14:14:18.3900	95	98	2	2	69	Kernel Started
01/15/2005	14:14:18.4200	95	98	3	2	10004	Valid Activation Key Detected :
01/15/2005	14:14:18.5350	98	170	4	2	52501	Target module loaded successfully
01/15/2005	14:14:18.6760	98	172	5	2	10004	Valid Activation Key Detected :
01/15/2005	14:14:18.9870	130	131	6	2	51501	Source module loaded successfully
01/15/2005	14:24:15.2070	130	132	7	2	72	Connection Request from ip://206.31.4.305
01/15/2005	14:24:16.3090	131	133	8	2	600002	Unified login provides ADMIN access
01/15/2005	14:24:40.9680	132	134	9	2	99	RepSet Modified: Exchange
01/15/2005	14:25:22.4070	134	131	10	2	71	Originator Attempting ip://206.31.4.305
01/15/2005	14:25:22.5030	134	131	11	2	0	Transmission Create to ip://206.31.4.305.
01/15/2005	14:25:22.6060	135	133	12	2	500000	Exchange is connected to ip://206.31.4.305
01/15/2005	14:25:23.5030	136	98	13	2	87	Start Replication on connection 1

The following table describes the information found in each column of the LogViewer.

Column #	Description
Column 1	The date the message was generated.
Column 2	The time the message was generated.
Column 3	The process ID.
Column 4	The thread ID.
Column 5	The sequence number is an incremental counter that assigns a unique number to each message. Useful when filtering out messages and comparing multiple LogViewer outputs.
Column 6	The type or level of message displayed: 1—Warning or error message 2—Informational message
Column 7	The LogViewer ID.
Column 8	The message text.

LogViewer log file

Each log file consists of a base name, a series number, and an extension.

- ◆ **Base Name**—The base name is determined by the application or process that is running:

Component	Log File Base Name
Storage Mirroring	dtlog
Management Console	mc
Failover Control Center	fcc
Text Client	dttxt
Command Line Client	dtcl

- ◆ **Series Number**—The series number ranges from 1 to 999. For example, Storage Mirroring begins logging messages to dtlog1. When this file reaches its maximum size, the next log file will be written to dtlog2. As long as log messages continue to be written, files dtlog3, dtlog4, dtlog5 will be opened and filled. When the maximum number of files is reached, which by default is 5, the oldest file is deleted when the sixth file is created. For example, when dtlog6 is created, dtlog1 is deleted and when dtlog7 is created, dtlog2 is deleted. When file dtlog999 is created and filled, dtlog1 will be recreated and Storage Mirroring will continue writing log messages to that file. In the event that a file cannot be removed, its number will be kept in the list, and on each successive file remove, the log writer will attempt to remove the oldest file in the list.
- ◆ **Extension**—The extension for each log file is .dtl.

Component	Sample Log File Names
Storage Mirroring	dtlog1.dtl, dtlog2.dtl
Management Console	mc1.dtl, mc2.dtl
Failover Control Center	fcc1.dtl, fcc2.dtl
Text Client	dttxt1.dtl, dttxt2.dtl
Command Line Client	dtcl1.dtl, dtcl2.dtl

Modifying the log file settings through the Management Console

1. To modify the maximum file size and the number of Storage Mirroring log files that are maintained, access the Server Properties dialog box by right-clicking a machine name in the left pane of the Management Console and selecting **Properties**.
2. Select the **Logging** tab.
3. At the top of the window, **Folder** indicates the directory where the log files are located. The default is the directory where the Storage Mirroring program files are installed.
4. Modify any of the options under **Messages and Alerts**, if necessary.
 - ◆ **Filename**—This field is no longer functional.
 - ◆ **Maximum Length**—Specify the maximum length of the log file. The default size is 1048576 bytes and is limited by the available hard drive space.
 - ◆ **Maximum Files**—Specify the maximum number of log files that are maintained. The default is 5 and the maximum is 999.

NOTE: If you change the **Maximum Length** or **Maximum Files**, you must restart the Storage Mirroring service for the change to take effect.

5. Click **OK** to save the changes.

Modifying the log file settings through the Text Client

1. To determine the current location where the log file is stored, use the `get` command with the `LogDir` option.

Command	<code>GET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	<code>GET <option> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>get netport</code>♦ <code>get VerifyLogName</code>♦ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

2. To modify the verify log file name, use the `set` command with the `LogDir` option. The `LogDir` option is any valid directory.

Command	<code>SET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	<code>SET <option>=<value> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>set netport=1100</code>♦ <code>set VerifyLogName="HA server.log"</code>♦ <code>set DefaultAddress="129.12.34.3"</code>
Notes	<ul style="list-style-type: none">♦ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed.♦ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

3. To determine the current number indicating how many log files can exist, use the `get` command with the `MaxNumberOfLogFiles` option.
4. To modify the number indicating how many log files can exist, use the `set` command with the `MaxNumberOfLogFiles` option. The default is 5 and the maximum is 999.
5. To determine the current maximum size of each log file, use the `get` command with the `MaxLogFileSize` option.
6. To modify the maximum size of each log file, use the `set` command with the `MaxLogFileSize` option. The default size is 1048576 bytes (1 MB) and is limited by the available hard drive space.

NOTE: If you change the `MaxNumberOfLogFiles` or `MaxLogFileSize`, you must restart the Storage Mirroring service for the change to take effect.

Viewing the log files

Use the `LogViewer` command from the directory where Storage Mirroring is installed to view a Storage Mirroring log file.

Command	LOGVIEWER
Description	The Storage Mirroring logging utility that records alerts (processing notifications, warnings, and errors) that occur during Storage Mirroring processing.
Syntax	LOGVIEWER [-PATH <path>] [-TYPE <number>] [-INCLUDE <list>] [-EXCLUDE <list>] [-NODATE] [-NOTIME] [-NOPID] [-NOTID] [-NOSEQ] [-NOTYPE] [-NOID] [-HELP]
Options	<ul style="list-style-type: none">◆ PATH—Allows you to specify a path to the log file◆ path—Specify the full path to the log file◆ TYPE—Allows you to filter the messages that are displayed◆ number—Specify 1 to display warning and error messages or specify 2 to display warnings, errors, and notifications◆ INCLUDE—Only includes specified IDs. All other IDs will not be displayed in the output◆ EXCLUDE—Excludes specified IDs. Ignore the specified IDs and display all others◆ list—A comma-separated list of IDs or ID ranges that follows the INCLUDE and EXCLUDE switches. A space should separate the switch from the list but within the list, there should be no spaces. Ranges are specified with a begin and end number and separated with a dash (-).◆ NODATE—Does not display the date in the output◆ NOTIME—Does not display the time in the output◆ NOPID—Does not display the process ID in the output◆ NOTID—Does not display the thread ID in the output◆ NOSEQ—Does not display the sequence number in the output◆ NOTYPE—Does not display the message type number in the output◆ NOID—Does not display the LogViewer ID in the output◆ HELP—Displays the command options
Examples	<ul style="list-style-type: none">◆ <code>LogViewer -type 2</code>◆ <code>LogViewer -include 200</code>
Notes	The default setting is <code>-type 2</code> which displays both type 1 and 2 messages.

LogViewer Messages

The following table describes the standard Storage Mirroring alerts that may be displayed in the LogViewer. The ID appears in column 7 of the LogViewer, and the message appears in column 8 of the LogViewer.

NOTE: In the following table, *con_id* refers to the unique connection ID assigned to each connection between a source replication set and a target.

ID	Message	Description
0	N/A	There are several log messages with this ID#. See the description in the Message column in the LogViewer file.
7	<ul style="list-style-type: none">◆ Synchronous ioctl returned STATUS_PENDING◆ Failed to reset Replication Flags. Replication may not be performed correctly.	<ul style="list-style-type: none">◆ Communication with the Storage Mirroring driver is not being performed correctly. A reboot is required to guarantee replication and data integrity.◆ An error occurred between the Storage Mirroring driver and recent changes to the replication set. The possible resolutions are:<ul style="list-style-type: none">◆ Undo the changes to the replication set◆ Stop and restart Storage Mirroring◆ Reboot the server
69	Storage Mirroring kernel started on <i>server_name</i>	The Storage Mirroring service was started on the Storage Mirroring server specified.
70	Storage Mirroring kernel stopped	The Storage Mirroring service was stopped on a Storage Mirroring server.
71	Originator attempting ip://xxx.xxx.xxx.xxx	A source is requesting to connect a replication set to a target machine.
72	Connection request from ip://xxx.xxx.xxx.xxx	A target machine has received a source machine's request to connect a replication set to the target.
73	Connected to ip://xxx.xxx.xxx.xxx	A source machine has successfully connected a replication set to a target machine.
74	Connection paused with ip://xxx.xxx.xxx.xxx	A network connection between the source and the target exists and is available for data transmission, but data is being held in queue and is not being transmitted to the target. This happens because the target machine cannot write data to disk fast enough. Storage Mirroring will resolve this issue on its own by transmitting the data in queue when the target catches up.
75	Connection resumed with ip://xxx.xxx.xxx.xxx	The transmission of data from the source machine to the target machine has resumed.
76	Connection failed to ip://xxx.xxx.xxx.xxx	An attempt to establish a network connection between a source machine and target machine has failed. Check your network connections and verify that the target machine is still online.
77	Connection lost with IP address <i>address</i>	The network connection previously established between a source machine and target machine has been lost. Check your network connections and troubleshoot to see why the connection was lost.
78	Auto-disconnect threshold has been reached.	The Storage Mirroring queue has exceeded its limit, and the auto-disconnect process will disconnect the source and target connection. The auto-reconnect process will automatically reestablish the connection if the auto-reconnect feature is enabled. If the auto-reconnect feature is not enabled, you must first verify that the connection between the source and target has been broken, and then manually reestablish the connection in the Management Console.

ID	Message	Description
79	Memory freed to bring Storage Mirroring memory usage below the limit	Data in the source queue has been sent to the target machine, bringing the pagefile below its limit.
80	Trying to auto-retransmit to ip://xxx.xxx.xxx.xxx	Storage Mirroring is attempting to automatically reconnect previously established source and target connections after a server reboot or auto-disconnect. This is also referred to as the auto-reconnect process.
81	Schedule transmit start to target	A scheduled transmission of data from a source machine to a target machine has started. See the description in the Message column in the LogViewer file.
82	Schedule transmit end to target	A scheduled transmission of data from a source machine to a target machine has ended. See the description in the Message column in the LogViewer file.
85	repset has been auto-disconnected	Storage Mirroring automatically disconnects the source and target connection because the queue size has reached a specified size for this action.
87	Start replication on connection <i>con_id</i>	Data has started replicating from a source machine to a target machine.
88	Stop replication on connection <i>con_id</i>	Data has stopped replicating from a source machine to a target machine.
89	Mirror started <i>con_id</i>	Data is being mirrored from a source machine to a target machine.
90	Mirror stopped <i>con_id</i>	The process of mirroring data from a source machine to a target machine has stopped due to user intervention or an auto-disconnect. (This means the mirroring process was not completed.)
91	Mirror paused <i>con_id</i>	The process of mirroring data from a source machine to a target machine has paused because the target machine cannot write the data to disk fast enough. Storage Mirroring will resolve this issue on its own by transmitting the data in queue when the target catches up.
92	Mirror resumed <i>con_id</i>	The process of mirroring data from a source machine to a target machine has resumed.
93	Mirror ended <i>con_id</i>	The process of mirroring data from a source machine to a target machine has ended.
94	Verification started <i>con_id</i>	The verification process of confirming that the Storage Mirroring data on the target is identical to the data on the source has started.
95	Verification ended <i>con_id</i>	The verification process of confirming that the Storage Mirroring data on the target is identical to the data on the source has ended.
97	Restore started <i>con_id</i>	The restoration process of copying the up-to-date data from the target back to the original source machine has started.
98	Restore completed <i>con_id</i>	The restoration process of copying the up-to-date data from the target back to the original source machine has been completed.
99	RepSet Modified: <i>repset_name</i>	This message means that the specified replication set has been modified.
100	Failover condition has been met and user intervention is required	Storage Mirroring has determined that the source has failed, and requires manual intervention to start the failover process.
101	Failover in progress!!!	The conditions for failover to occur have been met, and the failover process has started.

ID	Message	Description
102	Target full!	The disk to which data is being written on the target is full. This issue may be resolved by deleting files on the target machine or by adding another disk.
801	Auto-disconnect has occurred on IP address <i>address</i> with connection <i>con_id</i> . Disconnected replication set name: <i>repset_name</i> .	Auto-disconnect has occurred for the specified connection. This is due to the source queue filling up because of a network or target failure or bottleneck.
10001	Activation key is not valid.	An invalid activation code was identified when the Storage Mirroring service was started.
10002	Evaluation period has expired.	The evaluation license has expired.
10003	Activation code violation with machine <i>machine_name</i>	Duplicate single-server activation codes are being used on the servers, and Storage Mirroring is disabled.
10004	Valid activation key detected	A valid activation code was identified when the Storage Mirroring service was started.
51001	Source module failed to load	The Storage Mirroring source module failed to load. Look at previous log messages to determine the reason. (Look for messages that indicate that either the activation code was invalid or the user-configurable source module was not set to load automatically at startup.) The source module may have been configured this way intentionally.
51501	Source module loaded successfully	The Storage Mirroring source module was loaded successfully.
51502	Source module already loaded	The Storage Mirroring source module was already loaded.
51503	Source module stopped	The Storage Mirroring source module stopped.
52000	<ul style="list-style-type: none"> ◆ The target has been paused due to manual intervention. ◆ The target has been resumed due to manual intervention. 	The target has been paused or resumed through user intervention.
52000	Unfinished Op error	<p>This error message contains various Microsoft API codes. The text <code>Code -<x> Internal <y></code> appears at the end of this message. The code value indicates why the operation failed, and the internal value indicates the type of operation that failed. The most common code values that appear in this error message are:</p> <ul style="list-style-type: none"> ◆ (5) Permission denied: The account running the Storage Mirroring service does not have permission to update the file specified. ◆ (32) Sharing violation: Another application is using a particular file that Storage Mirroring is trying to update. Storage Mirroring will wait and try to update the file later. ◆ (112) Disk full: The disk to which data is being written on the target is full. This issue may be resolved by deleting files on the target machine or by adding another disk.
52501	Target module loaded successfully	The Storage Mirroring target module was loaded successfully.
52502	Target module already loaded	The Storage Mirroring target module was already loaded.
52503	Target module stopped	The Storage Mirroring target module stopped.
53001	File was missing from target	The verification process confirms that the files on the target are identical to the files on the source. This message would only appear if the verification process showed that a file on the source was missing from the target.


ID	Message	Description
53003	Could not read <i>filename</i>	Storage Mirroring could not read a file on the source machine because the file may have been renamed or deleted. For example, temporary files show up in queue but do not show up during transmission. (No user action required.)
54000	Kernel started	The Storage Mirroring service was started.
54001	Failover module failed to load	The Storage Mirroring failover module failed to load. Look at previous log messages to determine the reason.
54503	Failover module stopped	The Storage Mirroring failover module stopped.
99001	Starting source module low memory processing	The source's queue is full, and the auto-disconnect process will disconnect the source and target connection. The auto-reconnect process will automatically reestablish the connection if the auto-reconnect feature is enabled. If the auto-reconnect feature is not enabled, you must first verify that the connection between the source and target has been broken, and then manually reestablish the connection in the Management Console.
99999	Application is terminating normally	The Storage Mirroring service is shutting down normally.
503010	Asynclctl for status thread 178 terminated, terminating the status thread	A Storage Mirroring process monitors the state of the Storage Mirroring driver. When the service is shut down, the driver is shut down, and this process is terminated. (No user action required.)
600002	<ul style="list-style-type: none"> ◆ Unified login provides ADMIN access ◆ User <i>user</i> has <i>level</i> access (<i>x</i>) 	<ul style="list-style-type: none"> ◆ Using the current NT login grants ADMIN access. ◆ The listed user has listed access level and access level ID.
700000	The source machine <i>source_machine</i> is not responding to a ping.	This occurs when all monitored IP addresses on the source machine stop responding to pings. Countdown to failover will begin at the first occurrence and will continue until the source machine responds or until failover occurs.
800000	<ul style="list-style-type: none"> ◆ Active Directory GetHostSpns function call failed ◆ Active Directory RemoveSpns function call failed ◆ Active Directory AddSpns function call failed 	<ul style="list-style-type: none"> ◆ Storage Mirroring failed to get the host SPN (Service Principal Name) from Active Directory. ◆ Storage Mirroring failed to remove an SPN from Active Directory. ◆ Storage Mirroring failed to add a host SPN to Active Directory.

SNMP

SNMP, Simple Network Management Protocol, is the Internet's standard for remote monitoring and management of hosts, routers and other nodes and devices on a network. Storage Mirroring provides an SNMP sub-agent that monitors Storage Mirroring and can be managed from an SNMP Management Console.

Storage Mirroring installs two components to work with SNMP:

- ◆ The sub-agent is a program that installs and runs on the same machine as Storage Mirroring and gathers statistics, data, and traps. The sub-agent forwards the information to the SNMP agent, which relays the information to the manager. The Storage Mirroring SNMP sub-agent is included in the Storage Mirroring installation program.
- ◆ A Storage Mirroring .mib file is placed on the administrator's machine so that the management console can interpret the data sent from the sub-agent. The Storage Mirroring .mib file is dt.mib and meets SNMP standards.



Configuring SNMP

The Storage Mirroring .mib file will need to be loaded into your SNMP Management Console. Depending on the type of console you are using, this process might include compiling the .mib file. Reference your SNMP Management Console documentation for additional information.

SNMP statistics

The following tables identify the Storage Mirroring statistics provided through SNMP.

Object Type	Statistic	Description
General	dtActFailCount	The number of activation code errors
	dtAutoDisCount	The number of auto-disconnects
	dtAutoReCount	The number of auto-reconnects
	dtCurrentMemoryUsage	Amount of memory allocated from the Storage Mirroring memory pool
	dtFailedMirrorCount	The number of operations that failed to mirror because they couldn't be read on the source
	dtFailedRepCount	The number of operations that failed to be replicated because they could not be read on the source
	dtMirBytesGenerated	The number of bytes that have been transmitted by the mirroring process
	dtMirOpsGenerated	The number of mirror operations (create, modify, or delete) that have been transmitted by the mirroring process
	dtRepBytesGenerated	The number of bytes that have been transmitted by the replication process
	dtRepOpsGenerated	The number of operations (create, modify, or delete) that have been transmitted by the replication process
	dtUpTime	Time in seconds since Storage Mirroring was last started
Security	dtFailedLoginCount	The number of unsuccessful logins
	dtLoginCount	The number of successful logins

Object Type	Statistic	Description
Source	dtSourceState	0 —Source is not running 1 —Source is running without the replication driver 2 —Source is running with the replication driver.
	dtOpsDroppedCount	The number of file operations that have failed and will not be retried
	dtRetryCount	The number of file operations that have been retried
Target	dtTargetState	0 —Target is not running 1 —Target is running
	dtconBytesInMirQueue	The number of bytes in the mirror queue
	dtconBytesInRepQueue	The number of bytes in the replication queue
Connection	dtconBytesRx	The total number of bytes received from the target
	dtconBytesTx	The total number of bytes transmitted to the target
	dtconConnectTime	The duration of time since the connection was first established
	dtconIpAddress	The IP address of the connected machine. If at the source, then the IP address of the target. If at the target, then the IP address of the source.
	dtConnectionCount	The number of active connections between machines
	dtconOpsInAckQueue	The number of operations (create, modify, or delete) waiting for verification acknowledgements from the target
	dtconOpsInCmdQueue	The number of operations (create, modify, or delete) in the retransmit queue on the source
	dtconOpsInMirQueue	The number of mirror operations (create, modify, or delete) in the queue
	dtconOpsInRepQueue	The number of replication operations (create, modify, or delete) in the queue
	dtconOpsRx	The total number of operations (create, modify, or delete) received from the target
	dtconOpsTx	The total number of operations (create, modify, or delete) transmitted to the target
	dtconResentOpCount	The number of operations that were resent because of acknowledgement errors
	dtconState	The state of the active connection 0 —None. This indicates a connection has not been established. Statistics are still available for the source and target machines. 1 —Active. This indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.) 2 —Paused. This indicates a connection that has been paused. 4 —Scheduled. This indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations, time frame limitations, and so on). 8 —Error. This indicates a connection that is not transmitting because something has gone wrong (for example, lost connection). Only the Scheduled and Error states can coexist. All other states are mutually exclusive. SNMP will display a dtconState of 12 when the connection is in both a scheduled and an error state because this is the sum of the two values (4 + 8).

SNMP traps

The following tables identify the Storage Mirroring traps provided through SNMP.

Object Type	Trap	Description
General	dttrapLicenseViolationOnNetwork	A Storage Mirroring serial number conflict was identified on the network
Kernel	dttrapKernelStarted	Storage Mirroring has started
	dttrapKernelStopped	Storage Mirroring has stopped
	dttrapMemoryLimitReached	The Storage Mirroring memory pool limit has been reached
	dttrapMemoryLimitRemedied	The memory pool usage is below the maximum limit specified
Source	dttrapLicenseViolationStartingSource	The source cannot be started due to a license violation
	dttrapSourceStarted	Storage Mirroring source component has started
	dttrapSourceStopped	Storage Mirroring source component has stopped
Target	dttrapTargetFull	The target is full
	dttrapTargetStarted	Storage Mirroring target component has started
	dttrapTargetStopped	Storage Mirroring target component has stopped
Mirroring	dttrapMirrorEnd	Mirroring has ended
	dttrapMirrorPause	Mirroring has paused
	dttrapMirrorResume	Mirroring has resumed
	dttrapMirrorStart	Mirroring has started
	dttrapMirrorStop	Mirroring has stopped
Replication	dttrapReplicationStart	Replication has started
	dttrapReplicationStop	Replication has stopped
	dttrapRepSetModified	Replication set has been modified
Connection	dttrapAutoDisconnectEndConnection	Auto-disconnect has intentionally dropped the connection
	dttrapAutoDisconnectPauseTransmission	Auto-disconnect requested that the source pause any operation (create, modify, or delete) sending
	dttrapAutoDisconnectShutdown	Auto-disconnect forced Storage Mirroring to shutdown
	dttrapAutoDisconnectWriteQueue	Auto-disconnect has forced the queue to be written to disk
	dttrapAutoReconnect	Auto-reconnect needs to make a new connection
	dttrapConnectionFailed	The source to target connection was not successful
	dttrapConnectionLost	The source to target connection has been disconnected
	dttrapConnectionPause	The source to target connection has paused
	dttrapConnectionRequested	The source has requested a connection to the target
	dttrapConnectionRequestReceived	The target has received a connection request from the source
	dttrapConnectionResume	The source to target connection has resumed
	dttrapConnectionSucceeded	The source to target connection has been established
	dttrapScheduledConnectEnd	A scheduled end connection has been reached and the connection has been disconnected
	dttrapScheduledConnectStart	A scheduled connection has been established

Object Type	Trap	Description
Verification	dttrapVerificationEnd	Verification has ended
	dttrapVerificationFailure	Verification has failed
	dttrapVerificationStart	Verification has started
Failover	dttrapFailoverConditionMet	Manual intervention is required because failover has detected a failed source machine
	dttrapFailoverInProgress	Failover is occurring
Restoration	dttrapRestoreComplete	Restoration is complete
	dttrapRestoreStarted	Restoration has started

The Storage Mirroring program settings are the user-definable options that control Storage Mirroring application settings and Storage Mirroring processing settings. This includes options like the name of the verification log file and the size of the Storage Mirroring memory pool.

Changing the Settings Through the Management Console

All of the Storage Mirroring server settings are located in the Server Properties dialog box. To access this window, right-click a machine in the left pane of the Management Console and select **Properties**. The Server Properties dialog box contains eight tabs with all of the Storage Mirroring server settings.

General options

The **General** tab contains basic settings for the selected machine including the machine's identity and its Storage Mirroring functionality.

The screenshot shows the 'Server Properties' dialog box with the 'General' tab selected. The dialog has a title bar with a close button. Below the title bar are five tabs: 'Target', 'Database', 'Logging', 'E-mail Notification', and 'General'. The 'General' tab is active, showing fields for 'Nickname' (indy), 'Machine' (indy), 'Parent Group' (Servers\Sales\Sales US), 'Addresses' (192.168.2.32 : 1100), 'Client Transmit Port' (1100), and 'Discovery Method' (Active Directory (RnR)). Below these fields is a section for 'Operating System' showing 'Windows 2000 Advanced Server Version 5.0 (Build 2195: Service Pack 4)' and a 'License' section showing 'HP OpenView Storage Mirroring - Advanced Edition 4.4.1.0.1316.S' and an 'Activation Code' (01a2b000304c056d). At the bottom are 'OK', 'Cancel', and 'Help' buttons.

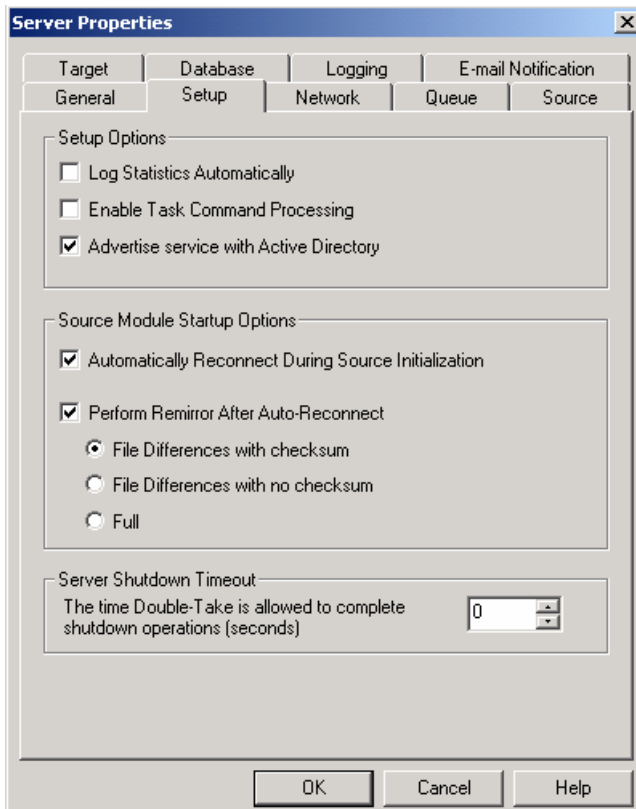
- ◆ **Nickname**—You can specify a machine nickname which is saved in the Management Console workspace. This nickname will only appear on the Management Console on this machine. It is not communicated across the network, but if this workspace is used on another machine, it will also appear there. For additional information on using a workspace on another machine, see [Storage Mirroring Workspaces](#) on page 2-12.
- ◆ **Machine**—This is the actual machine name. This field is not modifiable.
- ◆ **Parent Group**—The group in the Management Console server tree where the server is located. This field is not modifiable.
- ◆ **Addresses**—The IP address(es) for this machine are listed in this field. This information is not modifiable and is displayed for your information. The machine's primary address is listed first.
- ◆ **Client Transmit Port**—This field displays the port that the Management Console uses to send commands to a server. This port cannot be modified.
- ◆ **Discovery Method**—This field indicates the method in which the Management Console identifies the Storage Mirroring server.
 - ◆ **Manual**—A Storage Mirroring server was manually inserted into the Management Console server tree.
 - ◆ **Active Directory (RnR)**—A Storage Mirroring server is registered with Windows 200x Active Directory.
 - ◆ **Broadcast Heartbeat**—A Storage Mirroring server is broadcasting Storage Mirroring heartbeats.
- ◆ **Operating System**—Storage Mirroring identifies and displays the machine's operating system.
- ◆ **Storage Mirroring Version Information**—The Storage Mirroring version number and build number are displayed.

-
- ◆ **Activation Code**—This is the Storage Mirroring license which is required on every Storage Mirroring machine. There are three licenses available:
 - ◆ **Evaluation**—A license that has an end date built into the activation code. When the license expires, the software will no longer function. The same evaluation licenses can be used on multiple machines on a network.
 - ◆ **Single**—A license that is available on a per machine basis. Each machine is required to have a unique Storage Mirroring license whether it is functioning as a source, target, or both. A single license can only be used on one machine on a network.
 - ◆ **Site**—A license that is available to register every machine with the same Storage Mirroring license. This license is designed to be used on multiple machines on a network.

The activation code is a 16 character, alpha-numeric code. This code determines the Storage Mirroring license that is in use and can be changed without reinstalling Storage Mirroring if your Storage Mirroring license changes.

Setup options

The **Setup** tab consists of settings that are initialized when Storage Mirroring is started.



- ◆ **Log Statistics Automatically**—This check box indicates whether Storage Mirroring statistics logging, which logs to the `statistic.sts` file, will start when Storage Mirroring is started. For detailed information on the Storage Mirroring statistics logging, see [DTStat](#) on page 13-30.
- ◆ **Enable Task Command Processing**—This check box indicates whether task command can be processed on this server. If this setting is disabled and the server is a source, you can still submit tasks to be processed on the target, although task command processing must be enabled on the target. For more information, see [Processing task commands](#) on page 8-13.
- ◆ **Advertise service with Active Directory**—This check box indicates whether the Storage Mirroring service will register with Windows 200x Active Directory when the Storage Mirroring service is started. For more information, see [Windows 200x Active Directory](#) on page 15-7.
- ◆ **Automatically Reconnect During Source Initialization**—This check box indicates whether an auto-reconnect is to be performed if a source fails while a connection is established. For detailed information on the auto-reconnect feature, see [Auto-Disconnect and Auto-Reconnect](#) on page 6-7.
- ◆ **Perform Remirror After Auto-reconnect**—This check box indicates whether a remirror will be performed automatically after an auto-reconnect has occurred. For detailed information on the auto-remirror feature, see [Auto-Remirror](#) on page 7-4.
- ◆ **Server Shutdown Timeout**—The amount of time, in seconds, for the service to wait prior to completing the shutdown so that Storage Mirroring can persist data on the target in an attempt to avoid a remirror when the target comes back online. A timeout of zero (0) indicates waiting indefinitely and any other number indicates the number of seconds. The timeout setting only controls the service shutdown from the Storage Mirroring clients. It does not control the service shutdown through a reboot or from the Service Control Manager.

Network options

The **Network** tab consists of settings that control network communications. Any changes made to this tab will not take effect until the Storage Mirroring service has been restarted on the specified machine.

The screenshot shows the 'Server Properties' dialog box with the 'Network' tab selected. At the top, there are tabs for 'Target', 'Database', 'Logging', 'E-mail Notification', 'General', 'Setup', 'Network', 'Queue', and 'Source'. A warning icon and text state: 'Changing the values on this page will not take effect until the Storage Mirroring service has been restarted on the specified system.' Below this, the 'Interface' section contains: 'Default Address' (dropdown menu set to 'None'), 'Default Protocol' (dropdown menu set to 'TCP/IP'), 'Service Listen Port' (spin box set to 1100), 'Heartbeat Transmit Port' (spin box set to 1100), and 'Status Listen Port' (spin box set to 1105). The 'Heartbeat' section contains: 'Time between Heartbeats (seconds)' (spin box set to 3) and 'Missed Heartbeat Limit' (spin box set to 10). At the bottom are 'OK', 'Cancel', and 'Help' buttons.

◆ Interface

- ◆ **Default Address**—On a machine with multiple NICs, you can specify which address Storage Mirroring traffic will use. It can also be used on machines with multiple IP addresses on a single NIC.
- ◆ **Default Protocol**—The default protocol for all Storage Mirroring communications is the TCP/IP protocol. In the future, Storage Mirroring may support other communication protocols.
- ◆ **Service Listen Port**—Storage Mirroring servers use the Service Listen Port to send and receive commands and operations between two Storage Mirroring servers.
- ◆ **Heartbeat Transmit Port**—A Storage Mirroring server sends its heartbeats to the Heartbeat Transmit Port.
- ◆ **Status Listen Port**—Storage Mirroring servers use the Status Listen Port to listen for requests from the Management Console to update the at-a-glance monitoring information. For detailed information on at-a-glance monitoring, see [At-a-Glance monitoring](#) on page I3-6.

◆ Heartbeat

All Storage Mirroring servers, sources and targets, transmit a heartbeat. This heartbeat allows other Storage Mirroring machines, including Storage Mirroring clients, to locate and identify Storage Mirroring servers. The heartbeat is a broadcast UDP transmission. This heartbeat can be disabled, but Storage Mirroring will not auto-detect the Storage Mirroring machines to populate the Management Console.

- ◆ **Time Between Heartbeats**—This is the number of seconds between heartbeats. By default this is set to 3 seconds. If you set this option to 0, the heartbeats are disabled.
- ◆ **Missed Heartbeat Limit**—This is the number of heartbeats which can be missed before transmission is stopped and data is queued on the source.

Queue options

The **Queue** tab consists of settings for Storage Mirroring queue usage.

Server Properties

Target Database Logging E-mail Notification
General Setup Network Queue Source

⚠ Changing the queue location or memory usage will not take effect until the Storage Mirroring service has been restarted on the specified system.

Location
Folder: C:\Program Files\OpenView\Storage Mirroring ...
Free disk space (MB): 18829
ℹ A dedicated non-boot volume will offer the best performance and reliability.

Memory and Disk Usage
Maximum system memory for queue (MB): 128
Maximum disk space for queue (MB):
☒ Unlimited
Minimum Free Space (MB): 50

Queue Usage Alert Threshold
Alert at following queue usage percentage: 50

OK Cancel Help

- ◆ **Folder**—This is where the queue will be stored. For best results and reliability, you should select a dedicated, non-boot volume.
- ◆ **Maximum system memory for queue**—Storage Mirroring uses system memory to store data in queues. Specify the maximum amount of system memory to be used for the Storage Mirroring queues. The default is 128 MB. If you set it lower, Storage Mirroring will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Storage Mirroring will maximize system performance by not queuing to disk as soon, but the system may have to swap the memory to disk if the system memory is not available. In general, the amount of memory Storage Mirroring and other applications on the server are configured to use should be less than the amount of physical memory on the system to prevent low memory conditions.
- ◆ **Maximum disk space for queue**—This is the maximum amount of disk space in the specified **Folder** that can be used for Storage Mirroring queuing. When this limit is reached, Storage Mirroring will automatically begin the auto-disconnect process.
- ◆ **Minimum Free Space**—This is the minimum amount of disk space in the specified **Folder** that must be available at all times.

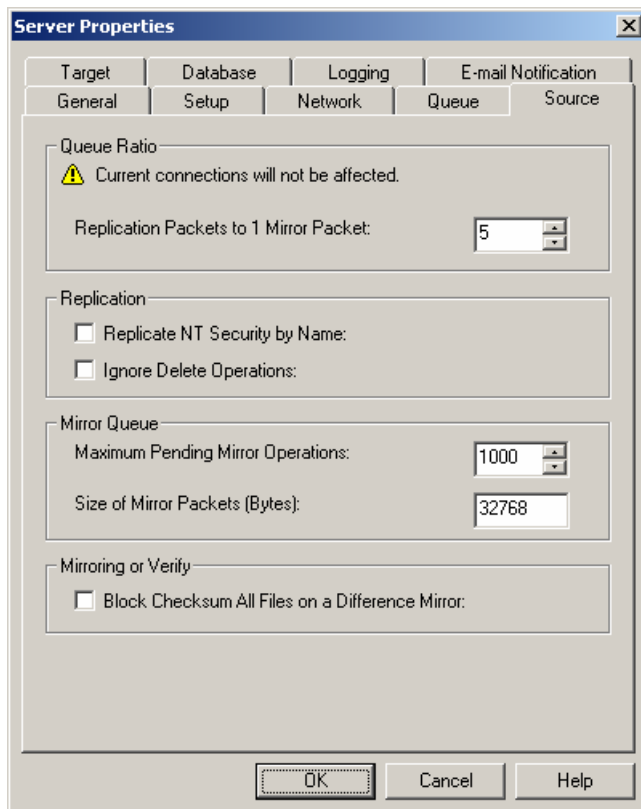
NOTE: The **Maximum disk space for queue** and **Minimum Free Space** settings work in conjunction with each other. For example, assume your queues are stored on a 10 GB disk with the **Maximum disk space for queue** set to 10 GB and the **Minimum Free Space** set to 500 MB. If another program uses 5 GB, Storage Mirroring will only be able to use 4.5 GB so that 500 MB remains free.

- ◆ **Alert at following queue usage percentage**—This is the percentage of the queue that must be in use to trigger an alert message.

NOTE: Scanning the Storage Mirroring queue files for viruses can cause unexpected results. If anti-virus software detects a virus in a queue file and deletes or moves it, data integrity on the target cannot be guaranteed. As long as you have your anti-virus software configured to detect the actual production data, the anti-virus software can clean, delete, or move an infected file and the clean, delete, or move will be replicated to the target. This will keep the target from becoming infected and will not impact the Storage Mirroring queues.

Source options

The **Source** tab consists of settings specific to the source module of Storage Mirroring.



- ◆ **Replication Packets to 1 Mirror Packet**—This option allows you to specify the ratio of replication packets to mirror packets that are placed in the source queue. Specify a larger number if you have a busy network that has heavy replication. Also, if you anticipate increased network activity during a mirror, increase this number so that the replication queue does not get too large.
- ◆ **Replicate NT Security by Name**—This option allows you to replicate permissions and attributes assigned to local (non-domain) users and groups. For detailed information, see [Windows Permissions](#) on page 8-2.
- ◆ **Ignore Delete Operations**—This option allows you to keep files on the target machine after they are deleted on the source. When a file is deleted on the source, that delete operation is ignored on the target. (All edits to files on the source are still replicated to the target; only deletions of whole files are ignored.)

NOTE: If a file is deleted using Windows Explorer or My Computer, the file is not actually deleted, from the file system perspective, but is moved to the Recycle Bin. Because Storage Mirroring sees this as a move to outside of the replication set and not as a delete, the file will still be deleted from the target even if you have **Ignore Delete Operations** selected.

If delete operations are ignored long enough, the potential exists for the target to run out of space. In that case, you can manually delete files from the target to free space for further replication.

◆ **Mirror Queue**

- ◆ **Maximum Pending Mirror Operations**—This option is the number of mirror operations that are queued on the source. The default setting is 1000. If, during mirroring, the mirror queued statistic regularly shows low numbers, for example, less than 50, this value can be increased to allow Storage Mirroring to queue more data for transfer.
- ◆ **Size of Mirror Packets**—The size of the mirror packets that Storage Mirroring transmits. The default setting is 32768 bytes.
- ◆ **Block Checksum All Files on a Difference Mirror**—This option allows a file difference mirror to check each block of data, regardless of the file attributes. If this option is not marked, Storage Mirroring will assume files are synchronized if their attributes match. For complete details on how this option interacts with the other mirror options, see [File Differences Mirror Options](#) on page 7-1.

NOTE: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the Block Checksum All option to ensure proper file comparisons.

Target options

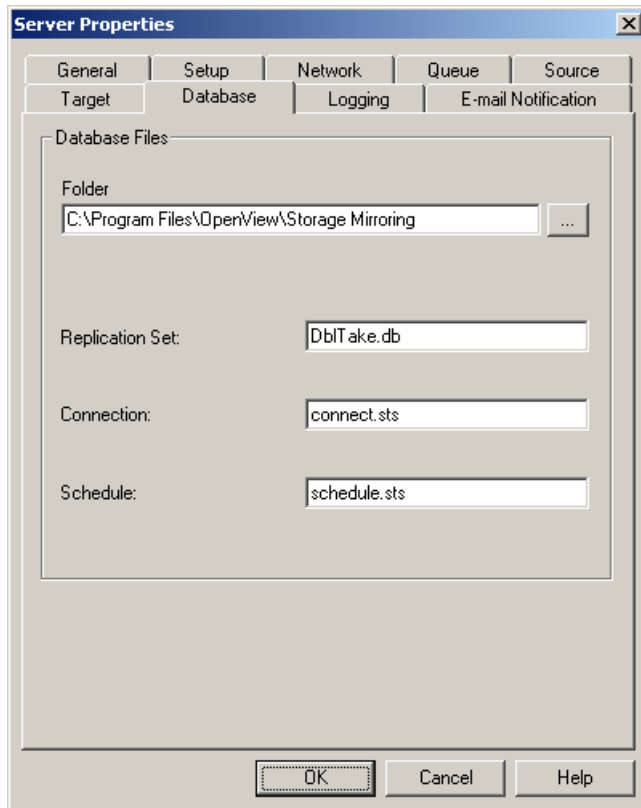
The **Target** tab consists of settings specific to the target module of Storage Mirroring.

The screenshot shows the 'Server Properties' dialog box with the 'Target' tab selected. The dialog has a title bar with a close button. Below the title bar are two rows of tabs: 'General', 'Setup', 'Network', 'Queue', 'Source' in the first row, and 'Target', 'Database', 'Logging', 'E-mail Notification' in the second row. The 'Target' tab is active. Inside the tab, there are two sections: 'Files' and 'Queues'. The 'Files' section contains two settings: 'Target Mirror Capacity High Percentage' with a value of 20, and 'Target Mirror Capacity Low Percentage' with a value of 15. The 'Queues' section contains one setting: 'Retry Delay for Incomplete Operations (seconds)' with a value of 3. At the bottom of the dialog are three buttons: 'OK', 'Cancel', and 'Help'.

- ◆ **Target Mirror Capacity High Percentage**—This option specifies the maximum percentage of Windows system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations. The default setting is 20.
- ◆ **Target Mirror Capacity Low Percentage**—This option specifies the minimum percentage of Windows system memory that can contain mirror data before the target signals the source to resume the sending of mirror operations. The default setting is 15.
- ◆ **Retry Delay for Incomplete Operations (seconds)**—This option specifies the amount of time, in seconds, before retrying a failed operation on the target. The default setting is 3.

Database options

The **Database** tab consists of settings that specify the name and location for the database files used by Storage Mirroring processes.



- ◆ **Folder**—Specify the directory where each of the database files on this tab are stored. The default is the directory where the Storage Mirroring program files are installed.
- ◆ **Replication Set**—This database file maintains which replication sets have been created on that machine along with their names, rules, and so on. The default file name is `Db1Take.db`.
- ◆ **Connection**—This database file maintains the active source/target connection information. The default file name is `connect.sts`.
- ◆ **Schedule**—This database file maintains any scheduling and transmission limiting options. The default file name is `schedule.sts`.

Logging options

The **Logging** tab consists of settings for various log files.

The screenshot shows the 'Server Properties' dialog box with the 'Logging' tab selected. The 'Folder' field is set to 'C:\Program Files\OpenView\Storage Mirroring'. The 'Messages & Alerts' section has 'Maximum Length (bytes)' set to 1048576 and 'Maximum Files' set to 5. The 'Verification' section has 'Filename' set to 'DtVerify.log', 'Maximum Length (bytes)' set to 1048576, the 'Append' checkbox checked, and 'Language' set to 'English'. The 'Statistics' section has 'Filename' set to 'statistic.sts', 'Maximum Length (bytes)' set to 10485760, and 'Write Interval (minutes)' set to 5. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

- ◆ **Folder**—Specify the directory where each of the log files on this tab are stored. The default is the directory where the Storage Mirroring program files are installed.
- ◆ **Messages & Alerts**
 - ◆ **Maximum Length**—Specify the maximum length of the Storage Mirroring alert log file. The default size is 1048576 bytes and is limited by the available hard drive space.
 - ◆ **Maximum Files**—Specify the maximum number of Storage Mirroring alert log files that are maintained. The default is 5, but the maximum is 999.

For additional information on the LogViewer utility, which displays the Storage Mirroring alerts, see [LogViewer](#) on page 13-39.

- ◆ **Verification**
 - ◆ **Filename**—The verification log is created during the verification process and details which files were verified as well as the files that are not synchronized. The default file name is `DtVerify.Log`.
 - ◆ **Maximum Length**—Specify the maximum length of the verification log file. The default maximum length is 1048576 bytes (1 MB).
 - ◆ **Append**—Mark the **Append** check box if you want to append each verification process to the same log file. If this check box is not marked, each verification process that is logged will overwrite the process log file. By default, this check box is selected.

For additional information on verification, see [Verification](#) on page 9-1.

- ◆ **Statistics**
 - ◆ **Filename**—The statistics log maintains connection statistics such as mirror bytes in queue or replication bytes sent. The default file name is `statistic.sts`. This file is a binary file that is read by the DTStat utility.
 - ◆ **Maximum Length**—Specify the maximum length of the statistics log file. The default maximum length is 10485760 bytes (10 MB). Once this maximum has been reached, Storage Mirroring begins overwriting the oldest data in the file.
 - ◆ **Write Interval**—Specify the frequency in which Storage Mirroring writes the statistical data to the statistics log file. The default is every 5 minutes.

For detailed information, see [DTStat](#) on page 13-30.

Notification options

The **E-mail Notification** tab allows you to specify e-mail notification settings. When enabled, the messages logged to the Windows Event Viewer will be sent as e-mail messages to the addresses specified. The subject of the e-mail will contain an optional prefix, the server name where the Event Viewer message was logged, the Event Viewer ID, and the severity level (information, warning, or error). The text of the Event Viewer message will be displayed in the body of the e-mail message.

In order to enable e-mail notification, the server must have Internet Explorer 5.0 or later. Additionally, if you are using Windows NT 4.0, you must also have Windows Management Instrumentation (WMI) installed.

Select or deselect **Enable notification** to turn this feature on or off. Any notification settings specified are retained when **Enable notification** is disabled.

The screenshot shows the 'Server Properties' dialog box with the 'E-mail Notification' tab selected. The 'Enable notification' checkbox is checked. The 'E-mail Settings' section includes fields for 'Mail Server (SMTP)' (mail.domain.com), 'Log on to SMTP server' (checked), 'Username' (administrator), 'Password' (masked), 'From Address' (HPQVSMAdmin@domain.com), 'Send To' (a list of email addresses: admin@domain.com, administrator@domain.com, IT_Support@domain.com), and 'Subject Prefix' (Storage Mirroring Notification). The 'Add event description to subject' checkbox is also checked. The 'Filter Contents' section has 'Include' checkboxes for 'Information' (unchecked), 'Warning' (checked), and 'Error' (checked). There is an 'Exclude these Event IDs' field with an example: 4000,4002-4010. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

◆ E-mail Settings

- ◆ **Mail Server (SMTP)**—Specify the name of your SMTP mail server.
- ◆ **Log on to SMTP Server**—If your SMTP server requires authentication, enable **Log on to SMTP Server** and specify the **Username** and **Password** to be used for authentication. Your SMTP server must support the LOGIN authentication method to use this feature. If your server supports a different authentication method or does not support authentication, you may need to add the Storage Mirroring server as an authorized host for relaying e-mail messages. This option is not necessary if you are sending exclusively to e-mail addresses that the SMTP server is responsible for.
- ◆ **From Address**—Specify the e-mail address that you want to appear in the From field of each Storage Mirroring e-mail message. The address is limited to 256 characters.
- ◆ **Send To**—Specify the e-mail address that each Storage Mirroring e-mail message should be sent to and click **Add**. The e-mail address will be inserted into the list of addresses. Each address is limited to 256 characters. You can add up to 256 e-mail addresses. If you want to remove an address from the list, highlight the address and click **Remove**. You can also select multiple addresses to remove by Ctrl-clicking.
- ◆ **Subject Prefix** and **Add event description to subject**—The subject of each e-mail notification will be in the format Subject Prefix : Server Name : Event Message Severity : Event Message ID : Event Message Description. The first and last components (Subject Prefix and Event Message Description) are optional. The subject line is limited to 150 characters.

If desired, enter unique text for the **Subject Prefix** which will be inserted at the front of the subject line for each Storage Mirroring e-mail message. This will help distinguish Storage Mirroring messages from other messages. This field is optional.

If desired, enable **Add event description to subject** to have the description of the Event Viewer message appended to the end of the subject line. This field is optional.

- ◆ **Filter Contents**—Specify which Event Viewer messages that you want to be sent via e-mail. Specify **Information**, **Warning**, and/or **Error**. You can also specify which Event Viewer messages to exclude based on the Event ID. Enter the Event IDs as a comma or semicolon separated list. You can indicate ranges within the list.

NOTE: You can test e-mail notification by specifying the options on the E-mail Notification tab and clicking **Test**. If desired, you can send the test message to a different e-mail address by selecting **Send To** and entering a comma or semicolon separated list of addresses. Modify the message text up to 1024 characters, if necessary. Click **Send** to test the e-mail notification. The results will be displayed in a message box. Click **OK** to close the message and click **Close** to return to the E-mail Notification tab.

E-mail notification will not function properly if the application or system event logs are full. It is recommended that:

- ◆ These event logs be less than 80% full.
- ◆ These event logs overwrite old events. Disable the option "Do not overwrite events."
- ◆ The maximum size of these event logs be increased. Set the maximum size of the event logs to at least 1024 KB.

These settings can be accessed through the Event Viewer by right-clicking the Application log and selecting **Properties**. Repeat this for the System log also.

If an error occurs while sending an e-mail, an Event message will be generated. This Event message will not trigger an e-mail. Subsequent e-mail errors will not generate additional Event messages. When an e-mail is sent successfully, an Event message will then be generated. If another e-mail fails, one Event message will again be generated. This is a cyclical process where one Event message will be generated for each group of failed e-mail messages, one for each group of successful e-mail messages, one for the next group of failed messages, and so on.

If you start and then immediately stop the Storage Mirroring service, you may not get e-mail notifications for the log entries that occur during startup.

Changing the Settings Through the Text Client

Changing Storage Mirroring settings using the Text Client can be confusing because the settings for a Storage Mirroring source or target are stored in the same location as the settings for the Storage Mirroring Text Client. Therefore, if you are running a Storage Mirroring source or target on the same machine that you are running the Text Client, and you update a client setting, you are also updating that setting on the source and/or target. The reverse is also true. If you update a setting on a source or target that is also running the Text Client, you are also updating that client setting.

All of the Storage Mirroring program settings are accessible using the `get` command. This command retrieves the current value of the specified setting.

Command	<code>GET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	<code>GET <option> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>get netport</code>♦ <code>get VerifyLogName</code>♦ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

Any setting can be changed using the `set` command. This command applies a new value to the setting.

Command	<code>SET</code>
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	<code>SET <option>=<value> [<i>machine</i>]</code>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>set netport=1100</code>♦ <code>set VerifyLogName="HA server.log"</code>♦ <code>set DefaultAddress="129.12.34.3"</code>
Notes	<ul style="list-style-type: none">♦ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to <code>ClientLogName</code> to take effect. And the service must be restarted to cause a change in the modules loaded if the <code>LoadSourceTarget</code> setting is changed.♦ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

Storage Mirroring Text Client program settings

The following table contains a complete listing of each Storage Mirroring program setting.

Program Setting	Description
ActivationCode	16-character Storage Mirroring activation code Values: Found on registration card, if not an evaluation copy Default: n/a GUI Setting: Server Properties, General tab, License, Activation Code
AdapterFlags	Specifies the adapter to use when establishing a connection Values: (1) Compression (2) Encryption (4) Network Data Representation Default: 4 GUI Setting: None Compression is no longer used. This option should not be changed.
Advertisement	Indicates how the Management Console populates the server tree Values: (1) Uses the Storage Mirroring heartbeat (4) Uses Windows 200x Active Directory (5) Uses both the Storage Mirroring heartbeat and Windows 200x Active Directory Default: 5 GUI Setting: Management Console, File, Options, Configuration tab, Automatic Service Discovery Changes to this setting require the Storage Mirroring service to be restarted for the change to take effect. If Active Directory Advertisement is enabled (set to 4 or 5), there is a 200 byte impact on the Active Directory service for each server that registers. The Storage Mirroring service registers with Active Directory at startup and unregisters at shutdown.
AllFailover	Specifies which IP addresses to failover Values: (0) Failover only monitored IP addresses (1) Failover all IP addresses Default: 1 GUI Setting: Failover Control Center, Monitor Settings, Items to Failover, IP Address(es)
AllMustFail	Specifies whether or not all IP addresses must fail for failover to take place Values: (0) any IP address can fail (1) All IP addresses must fail Default: 1 GUI Setting: Failover Control Center, Monitor Settings, Failover Trigger, All Monitored IP Addresses Fail
AutoReconnect	Specifies whether to reinstate the target connection(s) when the source machine is brought online after a source machine failure Values: (0) Do not reconnect (1) Reconnect Default: 1 GUI Setting: Server Properties, Setup tab, Source Module Startup Options, Automatically Reconnect During Source Initialization

Program Setting	Description
AutoRemirror	<p>Specifies whether to remirror when a source is brought online after an auto-disconnect</p> <p>Values: (0) Do not remirror (1) Perform a file differences checksum mirror (2) Perform a full mirror (3) Perform a file differences mirror</p> <p>Default: 1</p> <p>GUI Setting: Server Properties, Setup tab, Source Module Startup Options, Perform Remirror After Auto-Reconnect</p>
AutoRemirrorRetry	<p>Specifies how often, in seconds, the source should check for connections that have been reconnected but still need to be mirrored</p> <p>Values: any integer</p> <p>Default: 30</p> <p>GUI Setting: None</p>
AutoRetransmit	<p>Determines whether or not a source that has lost its connection with a target will attempt to reconnect to the target</p> <p>Values: (0) Do not attempt to reconnect (1) Attempt to reconnect</p> <p>Default: 1</p> <p>GUI Setting: None</p>
BackupDir	<p>Location on the target of the backup of the replication sets</p> <p>Values: any valid path</p> <p>Default: the location where the Storage Mirroring files were installed</p> <p>GUI Setting: None</p>
CalculateByVolume	<p>Calculates the approximate size of a replication set by using the size of the volume and subtracting the free space</p> <p>Values: (0) Disabled (1) Enabled</p> <p>Default: 0</p> <p>GUI Setting: None</p>
CalculateOnConnect	<p>Specifies whether or not the replication set size should be calculated on connection</p> <p>Values: (0) Do not calculate on connection (1) Calculate on connection</p> <p>Default: 1</p> <p>GUI Setting: Connection Manager, Mirroring tab, Calculate Replication Set size on connection</p>
CaseSensitiveRepSet Queries	<p>Specifies whether the replication set names, directories, and file names should be case sensitive</p> <p>Values: (0) Not case sensitive (1) Case sensitive</p> <p>Default: 0</p> <p>GUI Setting: None</p> <p>If you change this value, the source service must be stopped and restarted to implement the change.</p>

Program Setting	Description
ChecksumAll	<p>Setting to allow for the difference checksum option on mirror, verify, or restore to ignore the date, time, and size of the file and perform a checksum calculation on all files</p> <p>Values: (0) Checksum using date, time, size comparison (1) Checksum all files regardless of the date, time, or file size</p> <p>Default: 0</p> <p>GUI Setting: Server Properties, Source tab, Mirroring or Verify, Block Checksum All Files on a Difference Mirror</p>
Cleaner	<p>Specifies if a clean mirror will delete files on the target before mirroring</p> <p>Values:(0) Do not delete files before mirroring (1) Delete files before mirroring</p> <p>Default: 0</p> <p>This option is only valid if you are using one of the Storage Mirroring text clients, have this option enabled, and use the <code>clean</code> option.</p> <p>GUI Setting: None</p>
ClientLog	<p>Indicates whether commands entered in the Text Client are tracked</p> <p>Values: (0) Not tracked (1) Tracked</p> <p>Default: 0</p> <p>GUI Setting: None</p> <p>If you change this value, the Text Client must be closed and restarted to implement the change.</p>
ClientLogName	<p>Name of the log file that tracks commands entered in the Text Client</p> <p>Values: any valid file name</p> <p>Default: <code>Dttext.log</code></p> <p>GUI Setting: None</p> <p>If the name of the log file is changed, the Text Client must be closed and restarted to start logging to the new log file.</p>
ClusterDir	<p>Location of a MSCS installation, if it exists</p> <p>Values: any valid path</p> <p>Default: determined by the MSCS installation</p> <p>GUI Setting: None</p>
ConnectionDelay	<p>Length of time, in seconds, used by the Management Console to wait when logging in and making connections.</p> <p>Values: any valid number</p> <p>Default: 30</p> <p>GUI Setting: Management Console, File, Options, Configuration tab, Network, Communication Timeout</p>
ConnectionFile	<p>Name of the database file containing connection information</p> <p>Values: any valid file name</p> <p>Default: <code>connect.sts</code></p> <p>GUI Setting: Server Properties, Database tab, Database Files, Connection</p>
DataPath	<p>The location of the Storage Mirroring file attribute, replication set, connection, and schedule database files</p> <p>Values: any valid path</p> <p>Default: the location where the Storage Mirroring files were installed</p> <p>GUI Setting: Server Properties, Database tab, Database Files, Folder</p>

Program Setting	Description
DefaultAddress	<p>The default primary IP address in a multi-homed server</p> <p>Values: any valid IP address that will act as your primary IP address for connecting the source to the target</p> <p>Default: <null></p> <p>GUI Setting: Server Properties, Network tab, Interface, Default Address</p>
DefaultProtocol	<p>The default protocol</p> <p>Values: (1) IP protocol</p> <p>Default: 1</p> <p>GUI Setting: Server Properties, Network tab, Interface, Default Protocol</p> <p>TCP/IP is the only protocol currently supported.</p>
DirUNetPort	<p>Port used for directed UDP communications</p> <p>Values: 1025 - 65535</p> <p>Default: 1105</p> <p>GUI Setting: Server Properties, Network tab, Interface, Network, Status Listen Port Management Console, File, Options, Network, Status Transmit Port</p> <p>If you change this value, the source service must be stopped and restarted to implement the change.</p>
DisableAttribute Replication	<p>Specifies whether or not attributes (read-only, hidden, and so on) are replicated to the target</p> <p>Values: (0) Enable attribute replication (1) Disable attribute replication</p> <p>Default: 0</p> <p>GUI Setting: None</p>
DropOpOnHandle Error	<p>Determines if an additional attempt is made to access a file by a Microsoft API call if the Storage Mirroring call fails.</p> <p>Values: (0) When opening a file using the Storage Mirroring driver fails, attempt to open the file using the Microsoft Win32 API (1) When opening a file using the Storage Mirroring driver fails, skip the file and document it in the Storage Mirroring log</p> <p>Default: 1</p> <p>GUI Setting: None</p> <p>If the value is set to 0 and the Win32 call also fails, Storage Mirroring will skip the file and document it in the Storage Mirroring log</p>
DTSetupType	<p>Used by the Storage Mirroring installation program to maintain the installation settings for an upgrade</p>
EmailEnabled	<p>Specifies if e-mail notification is enabled</p> <p>Values: (0) E-mail notification is disabled (1) E-mail notification is enabled</p> <p>Default: 0</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Enable notification</p> <p>This is a read-only setting. If you change this setting using the DTCL SET command or via the registry editor, e-mail notification will not automatically start. You must use the GUI or the DTCL EMAIL ENABLE command to start e-mail notification.</p>

Program Setting	Description
EmailExcludelds	<p>Identifies the Windows Event Viewer messages that are excluded from e-mail notification.</p> <p>Values: Comma or semicolon separated list of Event Viewer IDs. You can indicate ranges within the list.</p> <p>Default: None</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Exclude these Event IDs</p> <p>You must disable and re-enable e-mail notification for a change to take effect.</p>
EmailFromAddress	<p>Specifies the e-mail address that will appear in the From field of Storage Mirroring generated e-mail messages.</p> <p>Values: Any valid e-mail address, up to 256 characters</p> <p>Default: None</p> <p>GUI Setting: Server Properties, E-mail Notification tab, From Address</p>
EmailInclude Categories	<p>Specifies which Event Viewer messages are sent via e-mail</p> <p>Values: (1) Error messages will be sent via e-mail (2) Warning messages will be sent via e-mail (3) Information messages will be sent via e-mail</p> <p>Default: 1, 2</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Include</p> <p>You must disable and re-enable e-mail notification for a change to take effect.</p>
EmailNotificationList	<p>Specifies the e-mail address(es) that will receive Storage Mirroring generated e-mail messages.</p> <p>Values: A comma separated list of valid e-mail addresses, up to 256 addresses. Each address is limited to 256 characters.</p> <p>Default: None</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Send To</p>
EmailPassword	<p>The password required for SMTP server authentication</p> <p>Values: Any valid password text</p> <p>Default: None</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Log on to SMTP Server, Password</p> <p>Since the password is encrypted for security, this entry cannot be displayed or changed through the registry or through the Text Client.</p>
EmailServer	<p>The name of the SMTP server for e-mail notification</p> <p>Values: Any valid server name text</p> <p>Default: None</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Mail Server (SMTP)</p>
EmailSmtpLogin	<p>Specifies if SMTP server authentication for e-mail notification is enabled or disabled</p> <p>Values: (0) SMTP authentication is disabled (1) SMTP authentication is enabled</p> <p>Default: 0</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Log on to SMTP Server</p> <p>Your SMTP server must support the LOGIN authentication method to use this feature. If your server supports a different authentication method or does not support authentication, you may need to add the Storage Mirroring server as an authorized host for relaying e-mail messages. This option is not necessary if you are sending exclusively to e-mail addresses that the SMTP server is responsible for.</p>

Program Setting	Description
EmailSubjectDesc	<p>Specifies if the Event Viewer message will be appended to the end of the subject line for e-mail notification</p> <p>Values: (0) Event Viewer message is not included in the subject line (1) Event Viewer message is included in the subject line</p> <p>Default: 1</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Add event description to subject</p>
EmailSubjectPrefix	<p>Specifies unique text which will be inserted at the front of the subject line for each Storage Mirroring generated e-mail message. This will help distinguish the Storage Mirroring messages from other messages.</p> <p>Values: Any valid text</p> <p>Default: Storage Mirroring Notification</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Subject Prefix</p>
EmailUsername	<p>The user ID required for SMTP server authentication</p> <p>Values: Any valid user ID text</p> <p>Default: None</p> <p>GUI Setting: Server Properties, E-mail Notification tab, Log on to SMTP Server, Username</p> <p>Since the username is encrypted for security, this entry cannot be displayed or changed through the registry or through the Text Client.</p>
EnableDHCP	<p>Indicates if Storage Mirroring DHCP support is enabled</p> <p>Values: (0) Disabled (1) Enabled</p> <p>Default: 1</p> <p>GUI Setting: None</p>
EnablePerformance Tracking	<p>This entry will be used in the future.</p>
EnableRoot Encryption	<p>Specifies if the top level folder's encryption will match between the source and target.</p> <p>Values: (0) Disabled—If the folders do not match, there will be a Storage Mirroring performance improvement (1) Enabled—If the folders match, there may be a decrease in Storage Mirroring performance</p> <p>Default: 1</p> <p>GUI Setting: None</p>
EnableTaskCmd Processing	<p>Queues tasks inline with replication data</p> <p>Values: (0) Disable task command processing (1) Enable task command processing</p> <p>Default: 0</p> <p>GUI Setting: Server Properties, Setup tab, Setup Options, Enable Task Command Processing</p>
ExtensionNumber	<p>Used by the Storage Mirroring log files.</p>
FailAllAddresses	<p>Specifies whether or not to failover all IP addresses if only one monitored IP address fails</p> <p>Values: (0) Do not failover all IP addresses (1) Failover all IP addresses</p> <p>Default: 1</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Items to Failover, IP Address(es), All</p>

Program Setting	Description
FailbackHostname	<p>Returns the host SPN (Service Principle Name) to its original setting on failback</p> <p>Values: (0) Disabled (1) Enabled</p> <p>Default: 0</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Active Directory, Failback Hostname</p> <p>If you are using Active Directory, this option should be enabled or you may experience problems with failback.</p>
FailoverHostname	<p>Automatically removes the host SPN (Service Principle Name) from Active Directory on the source</p> <p>Values: (0) Disabled (1) Enabled</p> <p>Default: 0</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Active Directory, Failover Hostname</p> <p>If you are using Active Directory, this option should be enabled or you may experience problems with failover.</p>
FailoverOnRoute Failure	<p>Determines if failover will occur when receiving a router message back from an IP address on the network</p> <p>Values: (0) Failover will not occur when receiving a destination host unreachable message (1) Failover will occur when receiving a destination host unreachable message</p> <p>Default: 1</p> <p>GUI Setting: None</p>
FCCHelpPath	<p>The path and file name to the Failover Control Center online help file</p> <p>Default: C:\Program Files\OpenView\Storage Mirroring\online help\fcc\wwhelp.htm</p> <p>GUI Setting: None</p> <p>Do not change this value or the online help will not work.</p>
FileAccessRetry	<p>The number of times a failed driver call will be retried by the service.</p> <p>Values: 1 - 65535</p> <p>Default: 10</p> <p>GUI Setting: None</p>
FileQueueSize	<p>When a mirror is started, one thread reads from the disk and builds the file queue. Another set of threads reads files off of the queue and sends them to the target. This setting is the maximum size of the queue in entries. If you had 100 files to be mirrored and this was set to 16 (the default value), the first thread would fill the queue to a maximum of 16 entries.</p> <p>Values: 1 - 65535</p> <p>Default: 16</p> <p>GUI Setting: None</p> <p>This value must be set prior to starting the mirror process. The higher the number, the more memory that is used.</p>
ForceReplaceOn Failover	<p>Specifies additional failover options</p> <p>Values: (0) Use standard failover add / replace settings with no additional settings (1) Replace the target server name with that of the source and add the source IP address (2) Add the source server name to the target and replace the target IP address (3) Replace the target server name with that of the source and replace the target IP address</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Failover Method</p>

Program Setting	Description
ForceVerifyOnMirror	Specifies if verification will be performed with every difference mirror Values: (0) Verification is not performed with every difference mirror (1) Verification is performed with every difference mirror Default: 0 GUI Setting: None
HBExternalRate	Number of seconds between heartbeats Values: 0 - 65535 Default: 3 GUI Setting: Server Properties, Network tab, Heartbeat, Time between Heartbeats HP recommends a value that is less than 10 (see HBTTTL). Zero (0) turns the heartbeats off.
HBInternalRate	This entry is no longer used.
HBLoopback	This entry is no longer used.
HBTrace	Specifies whether the heartbeat debugging information is generated Values: (0) not generated (1) Generated Default: 0 GUI Setting: Server Properties, Network tab, Heartbeat, Show Heartbeat Messages in Logger
HBTTTL	Number of seconds without receiving a heartbeat before a remote machine is considered unavailable Values: 0 - 65535 Default: 10 GUI Setting: Server Properties, Network tab, Heartbeat, Missed Heartbeat Limit
HeartbeatIgnoreIPs	Indicates an IP address to be ignored by the Management Console. An ignored IP address will not be displayed in the Management Console servers tree. Values: an IP address in the format x.x.x.x Default: <null> GUI Setting: None
HPQueueRatio	Ratio of replication packets to one mirror packet Values: 0 - 65535 Default: 5 A HPQueueRatio of 5 indicates 5 replication packets to 1 mirror packet. GUI Setting: Server Properties, Source tab, Queue Ration, Replication Packets to 1 Mirror Packet If you change this value, the source service must be stopped and restarted to implement the change.
IgnoreArchiveBit	Specifies if the archive bit is compared during verification Values: (0) Archive bit is compared during a verification (1) Archive bit is not compared during a verification Default: 1 GUI Setting: None

Program Setting	Description
IgnoreDeleteOps	<p>Specifies if file and directory delete operations will be replicated to the target</p> <p>Values:(0) Delete operations are replicated to the target (1) Delete operations are not replicated to the target</p> <p>Default:0</p> <p>GUI Setting: Server Properties, Source tab, Replication, Ignore Delete Operations</p>
IgnorePPPAddresses	<p>Identifies if Storage Mirroring will use PPP (Point-to-Point Protocol) or SLIP (Serial Line Internet Protocol) adapters</p> <p>Values:(0) Storage Mirroring will send out heartbeats across the PPP/SLIP adapter (1) Storage Mirroring will not send out heartbeats across the PPP/SLIP adapter</p> <p>Default: 1</p> <p>GUI Setting: None</p>
IgnoreThumbnail Streams	<p>Specifies if thumbnails will be replicated to the target.</p> <p>Values:(0) Storage Mirroring will mirror and replicate all data streams (1) Storage Mirroring will not mirror or replicate any data about the alternate data streams for thumbnail images. When comparing data for a verification or difference mirror, alternate data streams for thumbnails will not be reported as different.</p> <p>Default: 1</p> <p>GUI Setting: None</p> <p>If you change this value to 0, you must restart the Storage Mirroring service in order for the Storage Mirroring driver to begin sending all data stream information to the service. If you change this value to 1, you do not need to restart the service.</p>
InstallPath	<p>Path specified during the Storage Mirroring installation</p> <p>Values: any valid path</p> <p>Default: specified during installation</p> <p>GUI Setting: None</p>
InstallVersionInfo	<p>Installation number specified during the Storage Mirroring installation</p>
IntermediateQueue Limit	<p>Amount of memory that may be allocated to the intermediate queue by the system memory manager when MemoryAllocatorMode is set to mixed mode (2).</p> <p>Values: 512-4194304</p> <p>Default: 65536 bytes</p> <p>GUI Setting: None</p> <p>The Storage Mirroring service must be restarted in order for this change to take effect.</p>
IPFailover	<p>Specifies whether or not to failover the IP addresses during failover</p> <p>Values:(0) Do not failover IP addresses (1) Failover IP addresses</p> <p>Default: 1</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Items to Failover, IP Address(es)</p>
LanguagesAvailable	<p>Specifies the Storage Mirroring language support that has been installed. Do not modify this setting. If you need to add or remove language support, use the Storage Mirroring installation program.</p>
LanguageSelected	<p>Specifies the language of the verification log</p> <p>Values: Depends on LanguagesSupported</p> <p>Default: Language used during the installation</p> <p>GUI Setting: Server Properties, Logging, Verification, Language</p>
LanguagesSupported	<p>Specifies the available languages for the verification log. Do not modify this setting. If you need to add or remove language support, use the Storage Mirroring installation program.</p>

Program Setting	Description
LoadSourceTarget	<p>Specifies the functionality of the Storage Mirroring service</p> <p>Values:(0) Neither the source nor target modules are loaded (1) Only the source module is loaded (2) Only the target module is loaded (3) Both the source and target modules are loaded</p> <p>Default: 3</p> <p>GUI Setting: None</p> <p>This setting will not take effect until the Storage Mirroring service has been restarted.</p>
LogAllOrphans	<p>Specifies whether orphan files are logged to the Storage Mirroring log on the target</p> <p>Values:(0) Do not log the orphan files to the Storage Mirroring log on the target (1) Log the orphan files to the Storage Mirroring log on the target</p> <p>Default: 0</p> <p>GUI Setting: Connection Manager/Restoration Manager, Orphans tab, Log Orphaned Files to Target List</p>
LogDir	<p>The location of the Storage Mirroring messages/alerts, verification, and statistics log files</p> <p>Values: any valid path</p> <p>Default: the location where the Storage Mirroring files were installed</p> <p>GUI Setting: Server Properties, Logging tab, Folder</p>
LogFile	<p>The name of the Storage Mirroring messages/alerts log file</p> <p>Values: any valid file name</p> <p>Default: DTLog</p> <p>GUI Setting: None</p>
LogMessageLevel	<p>Specifies the types of messages logged to the.dtl files</p> <p>Values: (0) No messages will be logged (1) Only alert messages will be logged (2) Alert and release messages will be logged (3) Alert, release, and debug messages will be logged</p> <p>Default: 2</p> <p>GUI Setting: None</p>
LoopbackID	<p>This entry is no longer used.</p>
MaxChecksumBlocks	<p>Specifies the number of checksum values retrieved from the target</p> <p>Values: any integer</p> <p>Default: 32</p> <p>GUI Setting: None</p>
MaxConnections	<p>Number of network requests that can be processed simultaneously. Windows is limited to 5 simultaneous requests.</p> <p>Values: 0 - 65535</p> <p>Default: 5</p> <p>GUI Setting: None</p> <p>HP recommends that you not change this value.</p>

Program Setting	Description
MaxDriverPagefile Size	<p>Specifies the amount of disk space, in MB, to use for Storage Mirroring driver paging</p> <p>Values: 64-4095</p> <p>Default: 256</p> <p>GUI Setting: None</p> <p>The Storage Mirroring driver pagefile is located in the directory specified by QJournalDir.</p>
MaxLogFileSize	<p>Maximum size, in bytes, of any .dtl log file</p> <p>Values: limited by available disk space</p> <p>Default: 1048576</p> <p>GUI Setting: Server Properties, Logging tab, Messages & Alerts, Maximum Length</p>
MaxLogPathname	<p>The maximum length of a file name (the entire volume\directory\filename including slashes, spaces, periods, extensions, and so on) that will be displayed in the Storage Mirroring log file and the Windows Event Viewer. File names longer than the MaxDisplayablePath will be truncated and will be followed by an ellipsis (...).</p> <p>Values: 1-32760</p> <p>Default: 32760</p> <p>GUI Setting: None</p>
MaxNumberofLog Files	<p>Maximum number of .dtl log files that can exist at one time. When Storage Mirroring creates a new .dtl file, if this number is exceeded, the oldest .dtl file is deleted.</p> <p>Values: 1 - 999</p> <p>Default: 5</p> <p>GUI Setting: Server Properties, Logging tab, Messages & Alerts, Maximum Files</p>
MaxRemoveOrphans OpSize	<p>Determines whether or not Storage Mirroring will send over multiple remove orphan operations. Storage Mirroring will send over the operations if a directory has more files than this number.</p> <p>Values: 0 - 131072</p> <p>Default : 1000</p> <p>GUI Setting: None</p>
MaxRetry	<p>A generic, application wide setting specifying the number of retry attempts for processes such as creating sockets, starting the service, and so on</p> <p>Values: any integer</p> <p>Default: 5</p> <p>GUI Setting: None</p>
MaxWriteChunkSize	<p>Maximum merged op size (in bytes) used during replication</p> <p>Values: 1 - 131072</p> <p>Default: 65536 (64K)</p> <p>GUI Setting: None</p>
MCHelpPath	<p>The path and file name to the Management Console online help file</p> <p>Default: C:\Program Files\OpenView\Storage Mirroring\online help\mc\wwhelp.htm</p> <p>GUI Setting: None</p> <p>Do not change this value or the online help will not work.</p>

Program Setting	Description
MemoryAllocatorCallbackMode	<p>Determines what action is taken when the MemoryQueueToDiskThreshold is met</p> <p>Values: (0) Auto-disconnect—Auto-disconnect processing is initiated when the MemoryQueueToDiskThreshold has been met. Connections will be reestablished when auto-reconnect occurs.</p> <p>(1) Pause—The Storage Mirroring service stops pulling operations from the driver when the MemoryQueue ToDiskThreshold has been met. The target will pause the source. The service will resume pulling operations when the target tells the source to resume.</p> <p>(2) Queue—The source and target begin queuing operations to disk.</p> <p>Default: 2</p> <p>GUI Setting: None</p>
MemoryQueueToDiskThreshold	<p>A percentage of QmemoryBufferMax that will trigger queueing to disk.</p> <p>Values: any valid percentage</p> <p>Default: 100</p> <p>GUI Setting: None</p>
MinCompressionFile Size	<p>The minimum file size, in bytes, that will be compressed. Files smaller than this size will not be compressed.</p> <p>Values: any file file size</p> <p>Default: 1024</p> <p>GUI Setting: None</p>
MirrorChunkSize	<p>Block size (in bytes) used in the mirroring process</p> <p>Values: 1 - 1048576</p> <p>Default: 32768 (32K)</p> <p>GUI Setting: Server Properties, Source tab, Mirror Queue, Size of Mirror Packets</p> <p>A higher block size value gives you better throughput, but only to a certain point, then it starts using more memory (this has to do with the way memory is allocated and deallocated). A lower block size value produces slower throughput, but uses memory efficiently. The optimal value is 32768.</p>
MirrorEncryptedFiles	<p>Specifies if Windows 200x encrypted files are mirrored</p> <p>Values: (0) Encrypted files are not mirrored</p> <p>(1) Encrypted files are mirrored</p> <p>Default: 1</p> <p>GUI Setting: None</p>
MirrorOverwrite	<p>Determines if the mirror process overwrites existing files</p> <p>Values: (0) never overwrite</p> <p>(1) always overwrite</p> <p>(2) overwrite if older</p> <p>Default: 1</p> <p>GUI Setting: None</p>
MirrorPrompting	<p>This entry is no longer used.</p>
MirrorQueueLimit	<p>Maximum number of mirror operations that can be queued on the source machine</p> <p>Values: 1 - 65535</p> <p>Default: 1000</p> <p>GUI Setting: Server Properties, Sorce tab, Mirror Queue, Maximum Pending Mirror Operations</p>

Program Setting	Description
MirrorZeroKFiles	<p>Specifies whether or not empty files, zero byte files, are included in a mirror</p> <p>Values: (0) Zero byte files are skipped and not mirrored to the target (1) All files are mirrored to the target</p> <p>Default: 1</p> <p>GUI Setting: None</p> <p>If MirrorZeroKFiles is enabled (0), zero byte files are skipped during a full mirror, file differences mirror, and a verification with synchronization.</p> <p>Zero byte files that contain alternate data streams that are not empty, will still be skipped if MirrorZeroKFiles is enabled.</p>
MissedHeartbeats	<p>Specifies the number of heartbeats that can go unanswered by the source before failover occurs, when using UDP heartbeats to monitor for failover</p> <p>Values: 1 - 65535</p> <p>Default: 20</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Missed Packets</p>
MissedPackets	<p>Specifies the number of requests sent by the target that go unanswered by the source before failover occurs, when using ICMP pings to monitor for failover</p> <p>Values: 1 - 65535</p> <p>Default: 5</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Missed Packets</p>
MoveOrphanedFiles	<p>Specifies if orphaned files are deleted or moved to the directory specified by MoveOrphansDir</p> <p>Values: (1) Move (0) Delete</p> <p>Default: 0</p> <p>GUI Setting: Connection/Restoration Manager, Orphans tab, Move/Delete Orphan Files</p>
MoveOrphansDir	<p>Indicates the name of the directory where orphaned files will be moved if MoveOrphanedFiles=1</p> <p>Values: any valid path</p> <p>Default: the location where the Storage Mirroring files were installed</p> <p>GUI Setting: Connection/Restoration Manager, Orphans tab, Move Orphaned Files to following location</p>
NameFailover	<p>Specifies whether or not to failover machine names</p> <p>Values: (0) Do not failover machine names (1) Failover machine names</p> <p>Default: 1</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Items to Failover, Server Name</p>
NetPort	<p>Port connection for TCP communications</p> <p>Values: 1025 - 65535</p> <p>Default: 1100</p> <p>GUI Setting: Server Properties, Network tab, Interface, Service Listen Port</p> <p>If you change this value, the source service must be stopped and restarted to implement the change.</p>
NetworkRetry	<p>Specifies the interval, in seconds, at which Storage Mirroring will attempt to reconnect to the target</p> <p>Values: any positive number</p> <p>Default: 10</p> <p>GUI Setting: None</p>

Program Setting	Description
NetworkStatusInterval	This entry is no longer used.
NetworkTimeout	<p>The maximum length of time, in seconds, to wait on a network connection. If data is not received over a network connection within the specified time limit, the connection is closed. During idle periods, Storage Mirroring sends small amounts of keep-alive data at an interval 1/6 of the NetworkTimeout value to keep the socket from being inadvertently closed.</p> <p>Values: any integer</p> <p>Default: 120</p> <p>GUI Setting: None</p>
PingFrequency	<p>Specifies, in seconds, how often a ping is sent to the source from a monitoring target</p> <p>Values: 1 - 65535</p> <p>Default: 5</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Monitor Interval</p>
PlaceholderAdapters	<p>For Windows NT 4, the name of the NIC adapter</p> <p>Values: determined by the NIC installation</p> <p>Default: determined by the NIC installation</p> <p>GUI Setting: Failover Control Center, Settings, Placeholders</p>
PlaceholderCounts	<p>For Windows NT 4, the number of IP address placeholders</p> <p>Values: 0 - 50</p> <p>Default: determined by the Storage Mirroring installation</p> <p>GUI Setting: Failover Control Center, Settings, Placeholders</p>
PostFailbackScript	<p>Location on the target where the post-failback script is located</p> <p>Values: Any valid path</p> <p>Default: <null></p> <p>GUI Setting: Failover Control Center, Monitor Settings, Scripts, Target, Post-Failback</p>
PostFailoverScript	<p>Location on the target where the post-failover script is located</p> <p>Values: Any valid path</p> <p>Default: <null></p> <p>GUI Setting: Failover Control Center, Monitor Settings, Scripts, Target, Post-Failover</p>
PreFailbackScript	<p>Location on the target where the pre-failback script is located</p> <p>Values: Any valid path</p> <p>Default: <null></p> <p>GUI Setting: Failover Control Center, Monitor Settings, Scripts, Target, Pre-Failback</p>
PreFailbackWait	<p>Specifies whether or not to wait for the pre-failback script to complete before finishing a failback</p> <p>Values: (0) Do not wait (1) Wait</p> <p>Default: 0</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Scripts, Target, Pre-Failback, Delay failback until script completes</p>
PreFailoverScript	<p>Location on the target where the pre-failover script is located</p> <p>Values: Any valid path</p> <p>Default: <null></p> <p>GUI Setting: Failover Control Center, Monitor Settings, Scripts, Target, Pre-Failover</p>

Program Setting	Description
PreFailoverWait	<p>Specifies whether or not to wait for the pre-failover script to complete before finishing a failover</p> <p>Values: (0) Do not wait (1) Wait</p> <p>Default: 0</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Scripts, Target, Pre-Failover, Delay failover until script completes</p>
PriorityClass	<p>The priority level at which the service runs.</p> <p>Values: 3</p> <p>Default: 3</p> <p>GUI Setting: None</p>
ProductCode	<p>Used by the Storage Mirroring installation program to maintain the installation settings for an upgrade</p>
QJournalDir	<p>The location where the queue is stored.</p> <p>Values: any valid path</p> <p>Default: the location specified during the installation</p> <p>GUI Setting: Server Properties, Queue tab, Location, Folder</p> <p>For best results and reliability, you should select a dedicated, non-boot volume. The queue should be stored on a fixed, local NTFS volume.</p> <p>This location also stores the Storage Mirroring driver pagefile.</p>
QJournalFileSize	<p>The size, in MB, of each queuing transaction log file.</p> <p>Values: any valid file size, up to 4095 MB</p> <p>Default: 5</p> <p>GUI Setting: None</p>
QJournalFreeSpaceMin	<p>The minimum amount of disk space, in MB, in the specified <code>QJournalDir</code> that must be available at all times.</p> <p>Values: dependent on the amount of physical disk space available</p> <p>Default: 50</p> <p>GUI Setting: Server Properties, Queue tab, Memory and Disk Usage, Minimum Free Space</p> <p>The <code>QJournalFreeSpaceMin</code> should be less than the amount of physical disk space minus <code>QJournalSpaceMax</code>.</p>
QJournalPreload	<p>The number of operations being pulled from the disk queue at one time. Do not modify this setting.</p>
QJournalSpaceMax	<p>The maximum amount of disk space, in MB, in the specified <code>QJournalDir</code> that can be used for Storage Mirroring queuing. When this limit is reached, Storage Mirroring will automatically begin the auto-disconnect process.</p> <p>Values: dependent on the amount of physical disk space available</p> <p>Default: Unlimited (equivalent to 4,294,967,295 MB)</p> <p>GUI Setting: Server Properties, Queue tab, Memory and Disk Usage, Maximum disk space for queue</p> <p>The unlimited setting allows the disk queue usage to automatically expand whenever the available disk space expands. Setting this option to zero (0) disables disk queuing.</p>

Program Setting	Description
QLogWriteThrough	<p>Specifies if the disk queues are write-through mode</p> <p>Values: (0) Disk queues are not write-through mode (1) Disk queues are write-through mode</p> <p>Default: 0</p> <p>GUI Setting: None</p> <p>While write-through mode may decrease the frequency of auto-disconnects, it may also decrease the performance of the source server.</p>
QmemoryBufferMax	<p>The amount of Windows system memory, in MB, that, when exceeded, will trigger queuing to disk.</p> <p>Values: dependent on the amount of physical memory available; minimum of 32 MB</p> <p>Default: 128 MB</p> <p>GUI Setting: Server Properties, Queue tab, Memory and Disk Usage, Maximum system memory for queue</p>
QueryonQuorumFile	<p>identifies if the Storage Mirroring service will reopen closed files on the quorum drive</p> <p>Values: (0) The Storage Mirroring service will not attempt to reopen a closed file on the quorum drive to get security descriptors or last modified times. (1) The Storage Mirroring service will attempt to reopen a closed file on the quorum drive to get security descriptors or last modified times.</p> <p>Default: 0</p> <p>GUI Setting: None</p>
QueueSizeAlert Threshold	<p>The percentage of the queue that must be in use to trigger an alert message in the Windows Event Viewer.</p> <p>Values: any valid percentage</p> <p>Default: 50</p> <p>GUI Setting: Server Properties, Queue tab, Queue Usage Alert Threshold, Alert at following queue usage percentage</p>
Registered	This entry is no longer used.
RemoveAllOrphans	<p>Specifies if all orphan files will be removed or only those based on RemoveOrphanTime</p> <p>Values: (0) Remove orphans based on the entry RemoveOrphansTime (1) Remove all orphans</p> <p>Default: 1</p> <p>GUI Setting: Connection/Restoration Manager, Orphans tab, Remove All Orphans</p>
RemoveOrphansTime	<p>Specifies the amount of time that must be expired before an orphan file is removed</p> <p>Values: 1 - 131072 minutes</p> <p>Default: 60 minutes</p> <p>GUI Setting: Connection/Restoration Manager, Orphans tab, Remove Orphans not modified within the following time period</p>
RemoveSharesOn Disconnect	<p>Specifies if shares are removed on the target machine when a Storage Mirroring replication set is disconnected from a target or a source machine is manually shutdown by the administrator. (Shares are not removed if either the source or target machines fail.)</p> <p>Values: (0) Remove shares from the target (1) Do not remove shares from the target</p> <p>Default: 1</p> <p>GUI Setting: None</p>

Program Setting	Description
ReplaceTarget	Specifies whether or not to replace the target identity with the source identity during a failover Values: (0) Do not replace (1) Replace Default: 0 GUI Setting: Failover Control Center, Monitor Settings, Failover Method
ReplicateNTSecurityByName	Determines whether or not Storage Mirroring replicates permissions and attributes assigned to local (non-domain) users and groups Values: (0) Do not replicate by name (1) Replicate by name Default: 0 GUI Setting: Server Properties, Source tab, Replication, Replicate NT Security by Name
RepSetDBName	Name of the database that contains replication set information Values: any valid file name Default: DbITake.db GUI Setting: Server Properties, Database tab, Database Files, Replication Set
RestoreOverwrite	Determines if the restoration process overwrites existing files Values: (0) never overwrite (1) always overwrite (2) overwrite if older Default: 2 GUI Setting: Restoration Manager, Servers tab, Overwrite existing files during restore
RestorePrompting	This entry is no longer used.
SaveStatFile	Determines if the statistic.sts (statistics logging) file is saved or overwritten Values: (0) overwrite (1) saved as statistic-old.sts Default: 1 GUI Setting: None
ScheduleFile	Name of the database file that contains transmission scheduling information Values: any valid file name Default: Schedule.sts GUI Setting: Server Properties, Database tab, Database Files, Schedule If you change the name of the database file, the source service must be stopped and restarted to start logging to the new database file.
ScheduleInterval	The number of seconds to wait before checking the transmission schedules to see if transmission should be started or stopped Values: 1 - 3600 Default: 1 GUI Setting: None
ShareFailover	Specifies whether or not to failover shares Values: (0) Do not failover shares (1) Failover shares Default: 1 GUI Setting: Failover Control Center, Monitor Settings, Items to Failover, Share(s)

Program Setting	Description
ShareUpdateInterval	<p>Specifies how often, in minutes, the share file will be sent to the target</p> <p>Values: 1 - 65535</p> <p>Default: 60</p> <p>GUI Setting: None</p>
ShutdownTimeout	<p>The amount of time, in seconds, for the service to wait prior to completing the shutdown so that Storage Mirroring can persist data on the target in an attempt to avoid a remirror when the target comes back online</p> <p>Values: any valid number of seconds where 0 (zero) indicates waiting indefinitely and any other number indicates the number of seconds</p> <p>Default: 0</p> <p>GUI Setting: None</p> <p>This setting only controls the service shutdown from the Storage Mirroring clients. It does not control the service shutdown through a reboot or from the Service Control Manager.</p>
SkipCompressionFile Ext	<p>A period delimited list of file types that are not compressed, even if compression is enabled.</p> <p>Values: any period delimited list of file types</p> <p>Default: mp3.exe.wmv.wma.qt.mpg.mpeg.zip.jpg.jpeg.tiff.tar.rar.cab</p> <p>GUI Setting: None</p>
SourcePendingAcks	<p>The number of operations received by the target queue in which the source is waiting for a response</p> <p>Values: 100 - 20,000</p> <p>Default: 2000</p> <p>GUI Setting: None</p>
SourcePostFailback Script	<p>Path on the source where the post-failback script is located</p> <p>Values: Any valid path</p> <p>Default: <null></p> <p>GUI Setting: Failover Control Center, Monitor Settings, Scripts, Source, Post-Failback</p>
StatsDriverLogFlags	<p>Indicates which driver statistics are logged to the Storage Mirroring log</p> <p>Values: (1) State (2) Operations (4) Paging (8) Timing</p> <p>Default: 0 By default, no driver statistics are logged.</p> <p>GUI Setting: None</p> <p>Use the sum of various values to log multiple driver statistics. For example, a setting of 5 would log paging and state statistics. A setting of 7 would log paging, operations, and state statistics. A setting of 15 would log all driver statistics.</p>
StatsFileName	<p>Default file for logging statistics</p> <p>Values: any valid file name</p> <p>Default: statistic.sts</p> <p>GUI Setting: Server Properties, Logging tab, Statistics, Filename</p>
StatsLoggingOn	<p>Specifies if Storage Mirroring logs statistics at startup</p> <p>Values: (0) Stats logging does not start when Storage Mirroring starts (1) Stats logging starts when Storage Mirroring starts</p> <p>Default: 0</p> <p>GUI Setting: Server Properties, Setup tab, Setup Options, Log Statistics Automatically</p>

Program Setting	Description
StatsMaxFileSize	Maximum size for the <code>statistic.sts</code> file Values: limited by available disk space Default: 10485760 (10 MB) GUI Setting: Server Properties, Logging tab, Statistics, Maximum Length
StatsMaxObjects	This entry is no longer used.
StatsPort	Port used by DTStat to gather Storage Mirroring statistics Values: 1025 - 65535 Default: 1106 GUI Setting: None If you change this value, the source service must be stopped and restarted to implement the change.
StatsShmSize	This entry is no longer used.
StatsWriteInterval	Interval, in minutes, in which statistics are written to the <code>statistic.sts</code> file Values: 0 - 65535 Default: 5 GUI Setting: Server Properties, Logging tab, Statistics, Write Interval
SystemMemoryLimit	Set by the Storage Mirroring service, each time it is started, to record the amount of available memory.
TargetPaused	Internal setting that indicates if the target machine is paused. Do not modify this setting.
TargetPausedVirtual	Internal setting that indicates which target machines are paused. Do not modify this setting.
TCPBufferSize	Size of the TCP/IP buffer in bytes. Values: 4096-7500000 Default: 37500 GUI Setting: None <p>The default setting creates a TCP window that will accommodate most environments. In most environments, this value will not need to be adjusted. However, if your Storage Mirroring network has a long end-to-end route and the throughput is not where you would expect it to be, then adjusting this parameter may have beneficial results. This value is the bandwidth delay product, which is calculated using the bandwidth of the network (in bits/second) times the round trip time (in seconds) between the two ends. Some recommended settings to improve Storage Mirroring throughput performance are:</p> <ul style="list-style-type: none"> ◆ Gigabit LAN—The setting should be changed to 375000. ◆ WAN—The setting should be adjusted to 130000. <p>While the calculations are fairly straight forward, the values that have been suggested are not exact because they depend on round trip time. Some improvements could be gained by adjusting these values either higher or lower. The value suited for your environment can best be determined through trial and error testing.</p>
TempDir	Temporary directory used when replicating Windows 200x encrypted files. GUI Setting: None
TGBlockOnConnect	Blocks the target path for all connections, regardless of the source, so that the data cannot be modified Values: (0) Target paths are not blocked (1) Target paths are blocked Default: 0 GUI Setting: None Do not modify this setting without assistance from Technical Support.

Program Setting	Description
TGCloseDelay	<p>The length of time, in milliseconds, a file is held open on the target</p> <p>Values: 0 - 2000</p> <p>Default: 1000</p> <p>GUI Setting: None</p> <p>If disk caching on the target is disabled either manually or by default (for example, by default on disks that host Active Directory database files), the target system may be slow during a mirror. If so, decreasing this setting to 100, 10, and 0 will result in incremental improvements, with 0 returning the system performance to normal.</p>
TGDisableAttribute Replication	<p>Specifies whether or not the attributes compression, ACL, and file mask are written to the target during mirroring and replication</p> <p>Values: (0) Enable attribute replication (1) Disable attribute replication</p> <p>Default: 0</p> <p>GUI Setting: None</p>
TGExecutionRetry Limit	<p>The number of times an unfinished operation will be retried on the target before it is discarded. If this value is set to zero (0), an operation will never be discarded and will be retried on the target until it is applied.</p> <p>Values: 0 - 65536</p> <p>Default: 0</p> <p>GUI Setting: None</p>
TGFileAlloc	<p>Indicates that Storage Mirroring allocates an entire file on the first write of a mirror operation</p> <p>Values: (0) Disabled (1) Enabled</p> <p>Default: 1</p> <p>GUI Setting: None</p> <p>To help eliminate file fragmentation on the target server, Storage Mirroring should allocate the entire file first. With extremely large files, the file allocation may take a long time. Therefore, you may want to disable the file allocation. If you disable file allocation, you will have more fragmentation on the target disk.</p>
TGMirrorCapacity High	<p>Maximum percentage of system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations.</p> <p>Values: 2-75</p> <p>Default: 20</p> <p>GUI Setting: Server Properties, Target tab, Files, Target Mirror Capacity High Percentage</p>
TGMirrorCapacity Low	<p>Minimum percentage of system memory that can contain mirror data before the target signals the source to resume the sending of mirror operations.</p> <p>Values: 1-75</p> <p>Default: 15</p> <p>GUI Setting: Server Properties, Target tab, Files, Target Mirror Capacity Low Percentage</p> <p>The maximum value for TGMirrorCapacityLow is either 75 or TGMirrorCapacityHigh, whichever is lower.</p>
TGRetryLocked	<p>Minimum number of seconds to wait before retrying a failed operation on a target</p> <p>Values: 0-65536</p> <p>Default: 3</p> <p>GUI Setting: Server Properties, Target tab, Queues, Retry Delay for Incomplete Operations</p>
TGThreadCount	<p>This setting is no longer used</p>

Program Setting	Description
TGUnfinishedOpEvent	<p>Specifies whether or not unfinished operations on the target are logged to the Event Viewer</p> <p>Values: (0) Unfinished operation messages are not logged (1) Unfinished operation messages are logged</p> <p>Default: 1</p> <p>GUI Setting: None</p>
TGWriteCache	<p>Specifies whether or not Storage Mirroring uses the intermediate cache</p> <p>Values: (0) Bypass the intermediate cache and write directly to disks (1) Do not bypass the intermediate cache</p> <p>Default: 1</p> <p>GUI Setting: None</p>
TGWriteFailureBeforeNotification	<p>Specifies the number of times an operation will be retried on the target before a notification is sent to the source to update the Target Status field of the Management Console.</p> <p>Values: 0-1024</p> <p>Default: 10</p> <p>GUI Setting: None</p> <p>If you change the setting to 0, the notification will be disabled.</p> <p>Changing this option will only affect how the status is changed on the Target Status field of the Management Console. To solve the underlying issue of why the operations are failing will require investigation into the Storage Mirroring log files.</p>
UNetPort	<p>Port connection for UDP communications</p> <p>Values: 1025 - 65535</p> <p>Default: 1100</p> <p>GUI Setting: Server Properties, Network tab, Interface, Heartbeat Transmit Port Management Console, File, Options, Configuration tab, Automatic Service Discovery, Heartbeat Advertisement, Port</p> <p>If you change this value, the source service must be stopped and restarted to implement the change.</p>
UpdateInterval	<p>Interval, in seconds, at which the Failover Control Center updates the monitored machines display</p> <p>Values: 1 - 9999</p> <p>Default: 1</p> <p>GUI Setting: Failover Control Center, Settings, Refresh Rate</p>
UpgradeCode	<p>Used by the Storage Mirroring installation program to maintain the installation settings for an upgrade</p>
UseEventLog	<p>Specifies whether or not messages are logged to the Windows Event Viewer</p> <p>Values: (0) Do not log messages to the Event Viewer (1) Log messages to the Event Viewer</p> <p>Default: 1</p> <p>GUI Setting: None</p>
UserIntervention	<p>Specifies whether or not user intervention is required to initiate a failover</p> <p>Values: (0) User intervention is not required (1) User intervention is required</p> <p>Default: 1</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Manual Intervention</p>
UseScheduledPause	<p>Used by Storage Mirroring for internal schedule processing. Do not modify this setting.</p>

Program Setting	Description
UseShareFile	<p>Specifies whether to create and use a share file or to use the shares that are currently stored in the target memory</p> <p>Values: (0) Use the shares that are currently stored in the target memory (1) Create and use a file containing the share information</p> <p>Default: 1</p> <p>GUI Setting: Failover Control Center, Monitor Settings, Use .SHR Share Mapping File</p>
VerifyLogAppend	<p>Specifies whether the <code>DTVerify.log</code> file will be appended to or overwritten</p> <p>Values: (0) Overwrite (1) Append</p> <p>Default: 1</p> <p>GUI Setting: Server Properties, Logging tab, Verification, Append</p>
VerifyLogLimit	<p>Maximum size of the <code>DTVerify.log</code> file in bytes</p> <p>Values: limited by available hard drive space, up to 4 GB</p> <p>Default: 1048576 (1 MB)</p> <p>GUI Setting: Server Properties, Logging tab, Verification, Maximum Length</p>
VerifyLogName	<p>Name of the verification log file</p> <p>Values: any valid file name</p> <p>Default: <code>DTVerify.log</code></p> <p>GUI Setting: Server Properties, Logging tab, Verification, Filename</p>
VerifyRetryInterval	<p>The time, in minutes, between when one verification fails and a retry is scheduled to begin.</p> <p>Values: any valid number</p> <p>Default: 3</p> <p>GUI Setting: None</p>
VerifyRetryLimit	<p>The number of time a verification will be retried.</p> <p>Values: any valid number</p> <p>Default: 5</p> <p>GUI Setting: None</p>
VersionInfo	<p>The version of Storage Mirroring that was installed</p> <p>Value: determined by the version installed</p> <p>Default: determined by the installation</p> <p>GUI Setting: None</p>
WarningPings	<p>This entry is no longer used.</p>

To ensure protection of your data, Storage Mirroring offers multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine running Storage Mirroring. To gain access to a particular Storage Mirroring source or target, the user must provide a valid operating system user name and password and the specified user name must be a member of one of the Storage Mirroring security groups. Once a valid user name and password have been provided and the Storage Mirroring source or target has verified membership in one of the Storage Mirroring security groups, the user is granted appropriate access to the source or target and the corresponding features are enabled in the client. Access to Storage Mirroring is granted on one of the following three levels:

- ◆ **Administrator Access**—All Storage Mirroring features are available for that machine. For example, this access level includes creating replication sets and establishing Storage Mirroring connections.
- ◆ **Monitor Access**—Statistics can be viewed on that machine, but Storage Mirroring features are not available. For example, this access level does not allow the user to create or modify replication sets or create or modify Storage Mirroring connections.
- ◆ **No Access**—The machine appears in the Storage Mirroring Management Console and can be pinged from the Storage Mirroring Text Client, but no other access is available.

Security Access Levels

The following table identifies which key Storage Mirroring features are available depending on the security access granted.

Storage Mirroring Feature	Administrator Access	Monitor Access
Modify Replication Sets	✓	
View Replication Sets and Rules	✓	✓
Control Connections, Mirroring, Replication, Verification, Restoration, Failover	✓	
View Connection, Mirroring, Replication, Verification, Restoration Processing Statistics	✓	✓
View Storage Mirroring Program Settings	✓	✓
Modify Storage Mirroring Program Settings	✓	


Security Advantages and Considerations

Storage Mirroring security provides machine-based protection allowing the network administrator to specify the individuals that can access all of the Storage Mirroring features as well as those that only have access to the Storage Mirroring statistics. This security prevents unauthorized users from modifying critical Storage Mirroring configurations like the data included or excluded from a replication set, changing a one-to-one configuration to a one-to-many configuration by adding another connection, or initiating a mirror or stopping replication.

NOTE: Although Storage Mirroring passwords are encrypted when they are stored, Storage Mirroring security design does assume that any machine running the Storage Mirroring client application is protected from unauthorized access. If you are running the Storage Mirroring client and step away from your machine, you must protect your machine from unauthorized access.

How Storage Mirroring Security Works

1. When any Storage Mirroring client machine attempts to access a source or target machine running on Windows, it will attempt to automatically logon to the source or target using the three methods below.



- ◆ The security credentials of the user currently logged into the Storage Mirroring client machine are sent to the Storage Mirroring source or target machine. From the security credentials, the source or target machine determines if the user is a member of either of the Storage Mirroring security groups and if so, grants the appropriate level of access.
- ◆ The last valid set of credentials (credentials previously granting either Administrator or Monitor level access) used to access each machine is recorded in the registry of the client machine. If the logon attempt using the credentials of the user currently logged in fails, a set of credentials is retrieved from the registry and is sent to the Storage Mirroring source or target. The Storage Mirroring source or target checks the validity of the credentials and determines if the user is a member of one of the Storage Mirroring security groups and then grants the appropriate level of access.

NOTE: You can disable the feature that maintains the security credentials in the registry. See [Clearing Maintained Security Credentials](#) on page 15-3 for details.

- ◆ Each valid set of credentials (credentials previously granting either Administrator or Monitor level access) used by the Storage Mirroring client application is recorded in a memory-resident credentials buffer maintained by the Storage Mirroring client application. If the logon attempts using the credentials of the user currently logged in or those credentials stored in the registry fails, a set of credentials is retrieved from the Storage Mirroring client application's credentials buffer and is sent to the source or target. This process is repeated until a valid set of credentials is found or the credentials buffer is exhausted.

NOTE: The credentials buffer is cleared each time the Storage Mirroring client application is closed.


2. The Storage Mirroring client tries each of these three methods until a set of credentials granting Administrator access is found. If no credentials granting Administrator access are found, the Storage Mirroring client attempts to find a set of credentials granting Monitor access. If no credentials grant Monitor access, the user must manually logon to the Storage Mirroring source or target by providing a user name, password, and domain.

NOTE: If a user name exists both on the local machine and on the network, Windows first attempts to login to the machine with the local user name and password and ignores the domain. If this fails, it then tries to login with the network user name, password and domain.

Clearing Maintained Security Credentials

By default, Storage Mirroring maintains cached security credentials on the client machine as described in [How Storage Mirroring Security Works](#) on page 15-2. To remove the credentials, follow these steps:

1. To access the credentials security option, click **File, Options** and select the Security tab.



2. To remove the security credentials, click **Clear Cached Security Credentials**.
3. Click **OK**.

Storage Mirroring Service

By default, Storage Mirroring is configured to log on as the system account. If you want to select a specific account to run the service, use these instructions:

1. Open the Storage Mirroring service settings.
 - ◆ **Windows NT**—Open the Control Panel Services applet and double-click the Storage Mirroring service.
 - ◆ **Windows 200x**—Select **Start, Programs, Administrative Tools, Services** and double-click the Storage Mirroring service.
2. Select the **Log On** tab (if you are using Windows NT, this is under **Log On As**), select **This Account** and enter a valid account.
3. Enter the password for this account.
4. Click **OK** to save these settings.

NOTE: The selected account will need to have additional user rights granted to it. If you are using Windows NT, begin with the next step. If you are using Windows 200x, begin with step 14.

5. To add operating system access for a particular user, select **Start, Programs, Administrative Tools (Common), User Manager or User Manager for Domains**.
6. Select **Policies, User Rights**.
7. Verify that the **Show Advanced User Rights** check box at the bottom of the User Rights Policy dialog box is selected.
8. From the pull-down list labeled **Right**, select **Act as part of the operating system**. (This option will not be displayed if the **Show Advanced User Rights** check box at the bottom of the dialog box is not marked.)
9. To grant this access to the user account that will be starting the Storage Mirroring service, click **Add**.
10. Click **Show Users** and select the account that will be starting the Storage Mirroring service.
11. Click **Add** on the Add Users and Groups dialog box and click **OK**.
12. Click **OK** on the Users Rights Policy dialog box to return to the User Manager.
13. Exit **User Manager or User Manager for Domains**. This user is now configured to run the Storage Mirroring service.
14. If you are using Windows 200x, select **Start, Programs, Administrative Tools, Local Security Policy**.
15. Expand the **Local Policies** folder and highlight the **User Rights Assignment** folder.
16. Double-click the option **Act as part of operating system** on the right pane of the screen.
17. Add the user that you selected to run the Storage Mirroring service and click **OK**.
18. Exit the Local Security Settings dialog box. This user is now configured to run the Storage Mirroring service.

NOTE: If domain-level policy settings are defined (through **Domain Security Policy, Security Settings, Local Policies, User Rights Assignment**), they will override local policy settings.

Storage Mirroring Groups

Storage Mirroring requires two local groups on each source and target for security. During the installation process, these groups are automatically created. The groups are assigned specific case-sensitive names:

Double-Take Admin

Double-Take Monitors

The groups are installed differently depending on your operating system.

- ◆ **Windows 200x**—The local administrator and the domain administrator are automatically added to the **Double-Take Admin** group.

NOTE: For a Windows 200x member server, if an Active Directory user is granted access to the Active Directory **Double-Take Admin** or **Double-Take Monitors** groups, the user or domain group must also be granted access to the local Storage Mirroring groups. If you are administering a Windows 200x domain controller, the Active Directory group will provide sufficient access.

- ◆ **Windows NT**—Two factors determine what accounts are added to the **Double-Take Admin** group:
 - ◆ The type of machine you are logged on to (domain controller or member server) when you install Storage Mirroring
 - ◆ The account used to install Storage Mirroring (local user, local administrator, or domain administrator account)

The following table shows what accounts are added to the **Double-Take Admin** group when you install Storage Mirroring on Windows NT.

Account Used to Install Storage Mirroring	Machine Type	Accounts Added to Admin Group
Local User	Member Server	Local User Local Administrator
Local Administrator	Member Server	Local Administrator
Domain Administrator	Member Server	Local Administrator Domain Administrator
Local User	Domain Controller	Local User Domain Administrator
Domain Administrator (Local Administrator)	Domain Controller	Domain Administrator

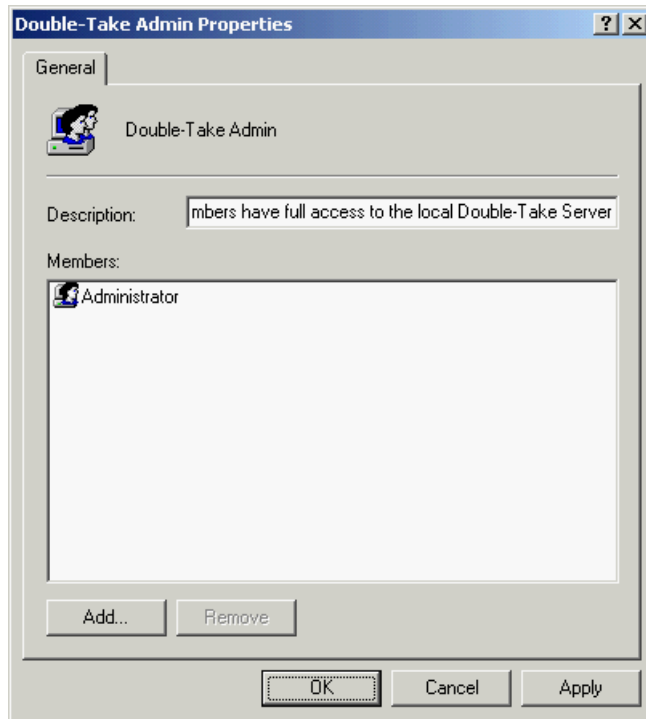
NOTE: For Windows NT, you should install Storage Mirroring as a domain administrator on either a member server or a domain controller (PDC/BDC) so that you will have administrator rights to all machines running Storage Mirroring.

Adding users to a group

Users that need administrator access to Storage Mirroring must be added to the **Double-Take Admin** group. All users that need monitor only access must be added to the **Double-Take Monitors** group. In both cases, local users, domain users, or global groups may be added to the local groups.

To add, delete, or modify users for a group, follow these steps:

1. Select **Start, Programs, Administrative Tools (Common), and User Manager**. (If you are on a domain controller, select **User Manager for Domains**.)
2. Double-click the group to be modified or highlight it and select **User, Properties**.



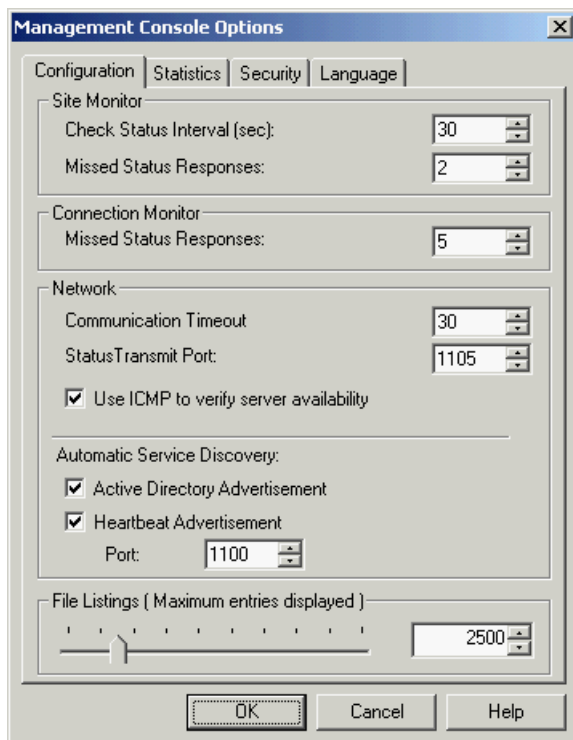
3. To add local users, domain users, and/or global groups to the group, click **Add**.
4. Select the local user, domain user, and/or global group to be included in the Storage Mirroring group.
5. Click **OK** to return to the Local Group Properties dialog box.
6. Click **OK** to return to the User Manager.
7. Exit the User Manager.

Windows 200x Active Directory

Active Directory, a central component of Windows 200x, manages information about the resources in a networking environment. The Storage Mirroring service automatically registers with Active Directory when the service starts. Storage Mirroring servers on the network can be located by the Storage Mirroring heartbeats, by Active Directory, or both. Active Directory strengthens network security by requiring the user account running the Storage Mirroring service to have specific privileges.

Disabling Active Directory

From the Management Console, select **File, Options** if you want to disable Active Directory.



Under Automatic Service Discovery, you can enable either **Active Directory Advertisement**, **Heartbeat Advertisement**, or both.

- ◆ **Active Directory Advertisement:** With this option enabled, only the Storage Mirroring servers registered in Active Directory will be displayed in the Management Console server tree.
- ◆ **Heartbeat Advertisement:** With this option enabled, all Storage Mirroring servers broadcasting Storage Mirroring heartbeats on the specified port number will be displayed in the Management Console server tree.

NOTE: By default, both **Active Directory Advertisement** and **Heartbeat Advertisement** are enabled. If both selections are enabled, **Active Directory Advertisement** takes precedence over **Heartbeat Advertisement**.

Configuring the Storage Mirroring service for Active Directory

The Storage Mirroring service must have privileges to modify Active Directory, if you want to use Active Directory registration. There are two options for assigning the privileges.

- ◆ **User account**—Assign a user account to the Storage Mirroring service and assign the Active Directory privileges to that user. To assign the user account and the privileges, follow the steps outlined in [Storage Mirroring Service](#) on page 15-4. Complete the first four steps to select the specific account to run the service, and then complete the remaining steps to assign the required privileges to the account. Refer to your Windows 200x reference guide for the specific privileges to assign.
- ◆ **Storage Mirroring Active Directory object**—Give the computer (or domain computers for all computers within a domain) read/write access to the Double-Take Instances object in Active Directory.
 1. Select **Start, Program, Administrative Tools, Active Directory Users and Computers**.
 2. Verify that **Advanced Features** is enabled on the **View** menu so that the **System** folder is displayed.
 3. Expand the **System** folder and select **WinsockServices**.
 - ◆ If you have not run the Storage Mirroring service under the domain administrator account or an account with update privileges for Active Directory, there will be no Storage Mirroring Active Directory instance listed. You will need to right-click on the **Winsock Services** folder to modify the setup for all Active Directory instances.
 - ◆ If you have run the Storage Mirroring service under the domain administrator account or an account with update privileges for Active Directory, **Double-Take Instances** will be listed. You can right-click **Double-Take Instances** to modify the Active Directory setup for the one instance or right-click on the **Winsock Services** folder to modify the setup for all Active Directory instances.
 4. Select the **Security** tab.
 5. Click **Add** and select the specific computer account you are running Storage Mirroring on or **Domain Computers** to allow all computers within the domain to update Active Directory.
 6. Grant both **Read** and **Write** access and click **OK**.

NOTE: If your corporate policies require that only the minimum required privileges be supplied, you can select only the permissions listed below by modifying the **Advanced** permissions for the account.

- ◆ List Contents
 - ◆ Read All Properties
 - ◆ Write All Properties
 - ◆ Read Permissions
-

A ► Storage Mirroring Commands

The Double-Take Command Language (DTCL) is a scripting language that can be used in either the Text Client or Command Line Client to manage and monitor Storage Mirroring components. It can also be used in script files to execute series and combinations of commands to meet specific needs. This appendix includes:

- ◆ **Conventions**—The conventions used in the command syntax
- ◆ **Storage Mirroring Commands**—An alphabetical list of all Storage Mirroring commands
- ◆ **DTCL Scripting**—Information on scripting with DTCL including variables, flow control, DTCL return values, and script examples.

Conventions

The following conventions are used to indicate command syntax for all Storage Mirroring related commands:

- ◆ UPPERCASE letters are used for commands that must be typed as shown.
- ◆ Underlined letters can be used as a shortcut for the command.
- ◆ *lowercase italicized* letters are variables such as file names, user names, or machine names. These items can also be substituted with defined variables as described in [DTCL Scripting](#) on page A-45.
- ◆ Angle brackets, < and >, surround required items that must be supplied with the command.
- ◆ Square brackets, [and], surround optional items that can be supplied with the command but are not required.
- ◆ The pipe character, |, separates items in a list
- ◆ Identifiers that contain a space or non-alphanumeric characters must be enclosed in quotation marks. For example:
 - ◆ “domain.com”
 - ◆ “machine name”
 - ◆ “129.65.35.45”

NOTE: You can specify an IP address or an IP address:port combination (separated by a colon) in place of a machine name in any DTCL command that allows for the entry of a machine name. For example, all of the following login commands are acceptable.

login mars administrator *****

login 216.234.244.47 administrator *****

login 216.234.244.47:1205 administrator *****

If you have modified your ports so that your source, target and/or clients are running on different ports and you are using DTCL scripts, then you will need to modify the scripts to include the port values as noted above. This only applies to source/target/client combinations that are communicating with each other on different ports. It does not apply between combinations of source/target pairs of servers that are using different ports.

Storage Mirroring Commands

The Storage Mirroring commands are listed in alphabetical order on the following pages. Each command includes the following:

- ◆ Command name
- ◆ Description of the command
- ◆ Exact syntax for using the command
- ◆ Options, if any
- ◆ Examples, if necessary
- ◆ Notes, if any

Compression Disable

Command	<code>COMPRESSION DISABLE</code>
Description	Disables compression
Syntax	<code>COMPRESSION DISABLE <con_id></code>
Options	<code>con_id</code> —Connection ID assigned to the source/target connection
Examples	<code>compression disable 1</code>

Compression Enable

Command	<code>COMPRESSION ENABLE</code>
Description	Enables compression
Syntax	<code>COMPRESSION ENABLE <con_id></code>
Options	<code>con_id</code> —Connection ID assigned to the source/target connection
Examples	<code>compression enable 1</code>

Compression List

Command	<code>COMPRESSION LIST</code>
Description	Identifies the compression level set and if compression is enabled
Syntax	<code>COMPRESSION LIST</code>

Compression Set

Command	COMPRESSION SET
Description	Sets the compression level
Syntax	COMPRESSION SET <con_id> <level>
Options	<ul style="list-style-type: none">◆ con_id—Connection ID assigned to the source/target connection◆ level—Any whole number from 1 to 3 where 1 is minimum compression and 3 is maximum compression
Examples	<pre>compression set 1 2</pre>
Notes	<ul style="list-style-type: none">◆ This command only sets the level of compression. It does not initiate compression.

ConID

Command	CONID
Description	Allows you to to assign the value of a connection ID to a variable or list the target and replication set for all connection IDs for a source
Syntax	<ul style="list-style-type: none">◆ <variable>=CONID <repset> TO <target>◆ CONID LIST <source_machine>
Options	<ul style="list-style-type: none">◆ variable—The name of the variable that you want to use to store the connecion ID◆ repset—The replication set that was used to establish the connection◆ target—The target that was used to establish the connection◆ source_machine—The name of the source server
Examples	<pre>\$con_id=conid exchange to jersey \$ConnectionNumber=conid sql to indy conid list jersey</pre>
Notes	<ul style="list-style-type: none">◆ The conid list and variable=conid commands are two separate commands.◆ If no machine name is specified in the conid list command, the active source will be used.

Connect

Command CONNECT

Description Establishes a connection between a replication set and a target machine

Syntax CONNECT <repset> TO <target_machine>
MAP EXACT
| MAP BASE <target_path>
| MAP <source_path> TO <target_path> [,...]
[MIRROR | NOMIRROR] [, REPLICATE | NOREPLICATE] [, MONITOR | NOMONITOR]
[, ORPHANS | NOORPHANS] [, COMPRESSION <level>] [CLEARRESTOREREQUIRED]

Options

- ◆ **repset**—Name of the replication set
- ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address
- ◆ **MAP EXACT**—Specifies that the replication set data will be sent to the same logical volume on the target (c:\data and d:\files is copied to c:\data and d:\files, respectively)
- ◆ **MAP BASE target_path**—Substitute a complete path, including the volume, for target_path and the data will be replicated to target_path\SrcVolName on the target machine
- ◆ **MAP source_path TO target_path**—Custom location that specifies each directory on the source and where that data will be copied to on the target machine
- ◆ **...**—Indicates that the source_path TO target_path option can be used more than once for each source directory in the replication set
- ◆ **MIRROR**—Automatically initiates a mirror when the connection is established
- ◆ **NOMIRROR**—Does not initiate a mirror when the connection is established
- ◆ **REPLICATE**—Automatically initiates replication when the connection is established
- ◆ **NOREPLICATE**—Does not initiate replication when the connection is established
- ◆ **MONITOR**—Specifies that the target is going to monitor the specified source machine for failover. The source machine must have already been defined as a monitor machine.
- ◆ **NOMONITOR**—Specifies that the target is not going to monitor the source machine for failover
- ◆ **ORPHANS**—Removes orphan files on the target
- ◆ **NOORPHANS**—Does not remove orphan files on the target
- ◆ **COMPRESSION level**—Sets the compression level for data being sent to the target using levels 1 through 3 with 1 being minimum compression and 3 being maximum compression. This option sets the level only. It does not enable compression.
- ◆ **CLEARRESTOREREQUIRED**—Clears the restore required flag and initiates the connection

Examples

- ◆ connect Exchange to jersey map exact
- ◆ connect sql to jersey map base d:\DTFiles\

Notes

- ◆ The default settings for this command are mirror, replicate, nomonitor, and noorphans.
- ◆ If a path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.
- ◆ If you are establishing a connection within a NAT environment, you will need to specify the port of the router after the IP address (seperated by a colon). See [NAT Environments](#) on page D-1 for more details.

Connect TDU

Command	CONNECT TDU
Description	Establishes a connection between a replication set and the <u>T</u> hroughput <u>D</u> iagnosics <u>U</u> tility to imitate a normal connection without transmitting any data across the network
Syntax	CONNECT <repset> TO TDU <file_name> [<i>connection flags</i>]
Options	<ul style="list-style-type: none">♦ <i>repset</i>—Name of the replication set♦ <i>file_name</i>—Name of the file to store the connection statistics generated by the TDU♦ <i>connection flags</i>—The same options available in the standard connect command
Examples	<ul style="list-style-type: none">♦ connect "Exchange Repset" to TDU♦ connect sql to TDU map c:\sql\data to e:\backup\sql\data
Notes	The statistic file that the TDU creates can be viewed using DTStat. By default, the file is called statistic.sts. To view the statistic file, type DTStat -f <filename>. For detailed information on DTStat, see DTStat on page 13-30.

Disconnect

Command	DISCONNECT
Description	Disconnects a specified source/target connection for the currently selected source.
Syntax	<u>DISCONNECT</u> <con_id *>
Options	<ul style="list-style-type: none">♦ <i>con_id</i>—Connection ID assigned to the source/target connection♦ *—Specifies all connection IDs.
Examples	<ul style="list-style-type: none">♦ disconnect 1♦ disconnect *

DTStat

Command	DTSTAT
Description	Starts the DTStats statistics logging utility from a command prompt.
Syntax	<pre>DTSTAT [-p][<i>-i</i> <interval>][<i>-t</i> <filename>] [<i>-f</i> <filename>][<i>-s</i> <filename>] [<i>-st</i> <filename>][<i>-IP</i> <address>][<i>-START</i> <mm/dd/yyyy hh:mm>][<i>-STOP</i> <mm/dd/yyyy hh:mm>] [<i>-SERVER</i> <ip_address> <port_number>]</pre>
Options	<ul style="list-style-type: none">◆ <i>-p</i>—Do not print the output to the screen◆ <i>-i</i> <i>interval</i>—Refresh from shared memory every <i>interval</i> seconds◆ <i>-t</i> <i>filename</i>—Save the data from memory to the specified binary file <i>filename</i>◆ <i>-f</i> <i>filename</i>—Reads from a previously saved binary file, <i>filename</i>, that was generated using the <i>-t</i> option instead of reading from memory◆ <i>-s</i> <i>filename</i>—Saves only the connection data from the data in memory to an ASCII, comma-delimited file, <i>filename</i>◆ <i>-st</i> <i>filename</i>—Saves only the target data from the data in memory to an ASCII, comma-delimited file, <i>filename</i>◆ <i>-f</i> <i>filename1</i> <i>-s</i> <i>filename2</i>—Saves only the connection data from a previously saved binary file, <i>filename1</i>, to an ASCII, comma-delimited file, <i>filename2</i>◆ <i>-f</i> <i>filename1</i> <i>-st</i> <i>filename2</i>—Saves only the target data from a previously saved binary file, <i>filename1</i>, to an ASCII, comma-delimited file, <i>filename2</i>◆ <i>-IP</i> <i>address</i>—Filters out the specified <i>address</i> in the IP address field and prints only those entries. Specify more than one IP address by separating them by a comma.◆ <i>-START</i> <i>mm/dd/yyyy hh:mm</i>—Filters out any data prior to the specified date and time◆ <i>-STOP</i> <i>mm/dd/yyyy hh:mm</i>—Filters out any data after the specified date and time◆ <i>SERVER</i> <i>ip_address</i> <i>port_number</i>—Connects DTStat to the specified IP address using the specified port number instead of to the local machine
Examples	<ul style="list-style-type: none">◆ <code>DTStat -i 30</code>◆ <code>DTStat -p -i 5 -t dtstat.log</code>◆ <code>DTStat -f dtstat.log -s dtstat.csv -start 02/02/2005 09:25</code>◆ <code>DTStat -server 206.31.4.51 1106</code>
Notes	<ul style="list-style-type: none">◆ This command is not case-sensitive.◆ If no options are specified, DTStat will print the output to the screen at an interval of every second.◆ If the statistics are not changing, DTStat will discontinue writing until statistics begin updating again.

Email Add

Command	EMAIL ADD
Description	Adds an e-mail address to the e-mail notification distribution list
Syntax	EMAIL ADD < <i>email_address</i> >
Options	<i>email_address</i> —A valid e-mail address enclosed in quotation marks
Examples	email add "administrator@domain.com"
Notes	<ul style="list-style-type: none">◆ You can repeat this command to add additional addresses to the distribution list.◆ You must have an active source specified for any e-mail command to work properly.

Email Disable

Command	EMAIL DISABLE
Description	Disables e-mail notification
Syntax	EMAIL DISABLE
Notes	You must have an active source specified for any e-mail command to work properly.

Email Enable

Command	EMAIL ENABLE
Description	Enables e-mail notification
Syntax	EMAIL ENABLE
Notes	You must have an active source specified for any e-mail command to work properly.

Email FromAddress

Command	EMAIL FROMADDRESS
Description	Specifies the e-mail address that will appear in the From field of Storage Mirroring generated e-mail messages
Syntax	EMAIL FROMADDRESS < <i>email_address</i> >
Options	<i>email_address</i> —A valid e-mail address enclosed in quotation marks
Examples	email fromaddress "Storage Mirroring@domain.com"
Notes	You must have an active source specified for any e-mail command to work properly.

Email Get

Command	EMAIL GET
Description	Displays e-mail notification configuration information
Syntax	EMAIL GET EMAILCONFIG FILTER
Options	<ul style="list-style-type: none">◆ EMAILCONFIG—Displays the current e-mail notification configuration◆ FILTER—Displays the filter configuration from the EMAIL SETFILTERINCLUDE switch
Notes	You must have an active source specified for any e-mail command to work properly.

Email MailServer

Command	EMAIL MAILSERVER
Description	Specifies the name of the SMTP mail server for e-mail notification
Syntax	EMAIL MAILSERVER <server_name ip_address> [username] [password]
Options	<ul style="list-style-type: none">◆ server_name—Specifies the name of the SMTP mail server◆ ip_address—Specifies the IP address of the SMTP mail server◆ username—The user ID required for SMTP server authentication◆ password—The password required for SMTP server authentication
Examples	<pre>email mailserver xchnng administrator *****</pre>
Notes	<ul style="list-style-type: none">◆ Your SMTP server must support the LOGIN authentication to supply a username and password. If your server supports a different authentication method or does not support authentication, you may need to add the Storage Mirroring server as an authorized host for relaying e-mail messages. This option is not necessary if you are sending exclusively to e-mail addresses that the SMTP server is responsible for.◆ You must have an active source specified for any e-mail command to work properly.

Email Remove

Command	EMAIL REMOVE
Description	Removes an e-mail address from the e-mail notification distribution list
Syntax	EMAIL REMOVE <email_address>
Options	email_address —An e-mail address, listed in the current distribution list, enclosed in quotation marks
Examples	<pre>email remove "administrator@domain.com"</pre>
Notes	You must have an active source specified for any e-mail command to work properly.

Email SetFilterInclude

Command	EMAIL SETFILTERINCLUDE
Description	Specifies which Event Viewer messages are sent via e-mail
Syntax	EMAIL SETFILTERINCLUDE [INFO, WARNING, ERROR] [EXCLUDEIDS "<ID1,ID2-ID4,...>"]
Options	<ul style="list-style-type: none">◆ INFO—Information messages will be sent via e-mail◆ WARNING—Warning messages will be sent via e-mail◆ ERROR—Error messages will be sent via e-mail◆ EXCLUDEIDS ID1,ID2-ID4,...—A comma separated list of IDs or ID ranges. A space should separate the EXCLUDEIDS switch from the list but within the list, there should be no spaces. Ranges are specified with a begin and end number and separated with a dash (-).
Examples	<ul style="list-style-type: none">◆ <code>email setfilterinclude warning, error</code>◆ <code>email setfilterinclude warning, error excludeids "4007,4012,4015"</code>◆ <code>email setfilterinclude excludeids "4000-4010"</code>
Notes	<ul style="list-style-type: none">◆ The default settings for this command are info and warning. No IDs are excluded, by default.◆ When changing the filter options, e-mail notification will automatically be stopped and restarted. If a failure occurs while stopping or restarting e-mail notification, a related error message will be output by DTCL.◆ You must have an active source specified for any e-mail command to work properly.

Email Subject

Command	EMAIL SUBJECT
Description	Specifies if additional text will be displayed in the subject of the e-mail message
Syntax	EMAIL SUBJECT [PREFIX <prefix> NOPREFIX] [DESCRIPTION NODESCRIPTION]
Options	<ul style="list-style-type: none">◆ PREFIX prefix—Text which will be inserted at the front of the subject line for each Storage Mirroring generated e-mail message. This will help distinguish the Storage Mirroring messages from other messages.◆ NOPREFIX—No prefix will be inserted at the front of the subject line for each Storage Mirroring generated e-mail message◆ DESCRIPTION—Append the message description to the end of the subject line◆ NODESCRIPTION—Do not append the message description to the end of the subject line
Examples	<ul style="list-style-type: none">◆ <code>email subject prefix "Storage Mirroring Notification" description</code>◆ <code>email subject prefix "Storage Mirroring Notification" nodescription</code>
Notes	<ul style="list-style-type: none">◆ If the prefix contains non-alphanumeric characters, it must be enclosed in quotation marks.◆ The entire subject line is limited to 150 characters.◆ You must have an active source specified for any e-mail command to work properly.

Email Test

Command	EMAIL TEST
Description	Sends a test message to demonstrate e-mail notification
Syntax	EMAIL TEST [SENDTO < <i>email_address</i> [, <i>email_address</i> , ...]>] [TEXT < <i>message_text</i> >]
Options	<ul style="list-style-type: none">♦ <i>email_address</i>—Specifies the e-mail address(es) to send the test message to if you do not want to use the e-mail addresses configured with the EMAIL ADD command. This is a comma separated list of addresses. The entire list should be enclosed in quotation marks. A space should separate the SENDTO switch from the list of addresses but within the list, there should be no spaces.♦ <i>message_text</i>—Text to be displayed in the body of the test e-mail message. The test message is limited to 1024 characters, if
Examples	<ul style="list-style-type: none">♦ email test sendto "administrator@domain.com,itsupport@domain.com"♦ email test sendto "admin@domain.com" text "This is a test message"
Notes	<ul style="list-style-type: none">♦ The default setting for this command are to use the addresses configured with the EMAIL ADD command. The default test message appends the word Test to the prefix defined in the EMAIL SUBJECT PREFIX command.♦ If the message contains non-alphanumeric characters, it must be enclosed in quotation marks.♦ You must have an active source specified for any e-mail command to work properly.

Environment

Command	ENVIRONMENT
Description	Displays a list of all Storage Mirroring machines available to the specified machine. Each machine is identified by machine name, IP addresses and whether or not the source and/or target modules are loaded. If no machine is specified, the information is provided for the machine currently specified as the source.
Syntax	ENVIRONMENT [<i>machine</i>]
Options	<i>machine</i> —Name of the machine to poll for environment information
Examples	environment indy

Exit

Command	EXIT
Description	Exits the Text Client and the Command Line Interactive client
Syntax	EXIT

Failback

Command	FAILBACK
Description	Initiates the failback process for the specified monitor machine
Syntax	FAILBACK <monitor> [ON <target_machine>] [REMONITOR NOREMONITOR]
Options	<ul style="list-style-type: none">♦ monitor—Name of the source machine designated as the monitor♦ target_machine—Name of the target machine♦ REMONITOR—Automatically continues monitoring the source machine after failback♦ NOREMONITOR—Automatically discontinues monitoring the source machine after failback.
Examples	failback indy on jersey
Notes	If you do not select <code>remonitor</code> or <code>noremonitor</code> , you will be prompted after failback is complete to select whether or not to continue monitoring.

Failover

Command	FAILOVER
Description	Manually initiates the failover process for the specified monitor machine
Syntax	FAILOVER <monitor> [ON <target_machine>]
Options	<ul style="list-style-type: none">♦ monitor—Name of the source machine designated as the monitor♦ target_machine—Name of the target machine
Examples	failover indy on jersey

Get

Command	GET
Description	The Text Client/DTCL sends a request to Storage Mirroring to return the value of a Storage Mirroring program setting. The setting is retrieved by Storage Mirroring from the registry.
Syntax	GET <option> [<i>machine</i>]
Options	<ul style="list-style-type: none">♦ option—See Storage Mirroring Text Client program settings on page 14-15.♦ machine—Name of the machine
Examples	<ul style="list-style-type: none">♦ <code>get netport</code>♦ <code>get VerifyLogName</code>♦ <code>get DefaultAddress</code>
Notes	If you do not specify a machine name, the value from the current source will be returned. If you have not identified an active source, no data will be returned.

Get Local

Command	GETLOCAL
Description	The Text Client/DTCL sends a request directly to the registry, bypassing Storage Mirroring, to return the value of a Storage Mirroring program setting.
Syntax	GETLOCAL < <i>option</i> >
Options	<i>option</i> —See <i>Storage Mirroring Text Client program settings</i> on page 14-15.
Examples	<ul style="list-style-type: none">◆ getlocal netport◆ getlocal VerifyLogName◆ getlocal DefaultAddress

GetEnvStr

Command	GETENVSTR
Description	Retrieves an operating system environment variable and stores the value in a variable.
Syntax	GETENVSTR < <i>env_variable</i> > \$< <i>variable_name</i> >
Options	<ul style="list-style-type: none">◆ <i>env_variable</i>—The name of the operating system environment variable that you want to store in the user defined variable◆ <i>variable_name</i>—The name of the variable you want to create. This variable will store the operating system environment variable specified.
Examples	getenvstr computename \$server
Notes	It is not necessary for either the environment or DTCL variable to exist when using this command. A non-existent environment variable will store a null string in the DTCL variable.

Help

Command	HELP
Description	Displays the DTCL commands and their syntax
Syntax	HELP
Notes	<ul style="list-style-type: none">◆ Press any key to scroll through the list of commands.◆ Press q to exit the help function.◆ You can also type dtcl help from the directory where the Storage Mirroring program files are installed to display the DTCL commands and their syntax.

Limit Bandwidth

Command	LIMIT BANDWIDTH
Description	Sets bandwidth restrictions for transmitting data from the source to the target
Syntax	LIMIT BANDWIDTH <bytes>, <seconds> TO <target_machine>
Options	<ul style="list-style-type: none">♦ bytes—Number of bytes to be transmitted♦ seconds—Maximum number of seconds to wait before transmitting again♦ target_machine—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>limit bandwidth 19300, 5 to jersey</code>
Notes	This command transmits in bursts, not bytes per seconds. The time identifies how long to wait before transmitting again. For example, if 5 seconds are specified and it only takes 2 seconds to send the specified bytes, Storage Mirroring will wait an additional 3 seconds before transmitting again.

Load Source

Command	LOAD SOURCE
Description	Loads the Storage Mirroring source module
Syntax	LOAD SOURCE <machine>
Options	machine —Name of the machine
Examples	<code>load source indy</code>

Load Target

Command	LOAD TARGET
Description	Loads the Storage Mirroring target module
Syntax	LOAD TARGET <machine>
Options	machine —Name of the machine
Examples	<code>load target jersey</code>

Login

Command	LOGIN
Description	Log on to a Storage Mirroring machine
Syntax	LOGIN < <i>machine</i> > < <i>username</i> > < <i>password</i> > [<i>domain</i>]
Options	<ul style="list-style-type: none">◆ <i>machine</i>—Name of the machine◆ <i>username</i>—Name of the user◆ <i>password</i>—Password associated with username.◆ <i>domain</i>—If logging in using a domain account, this is the domain name. If logging in using a local account, this is the machine name.
Examples	login indy administrator *****
Notes	<ul style="list-style-type: none">◆ The login command is not available when scrolling through the Text Client command history.◆ If characters in the password include non-alphanumeric characters, the password field must be enclosed in quotation marks.◆ The password cannot be a Storage Mirroring keyword. These are any DTCL command (source, target, and so on.) or any DTCL shortcut command (env, mon, rep, and so on).

Logout

Command	LOGOUT
Description	Logs off of a Storage Mirroring machine
Syntax	LOGOUT < <i>machine</i> >
Options	<i>machine</i> —Name of the machine
Examples	logout indy

LogViewer

Command	LOGVIEWER
Description	The Storage Mirroring logging utility that records alerts (processing notifications, warnings, and errors) that occur during Storage Mirroring processing.
Syntax	LOGVIEWER [-PATH <path>] [-TYPE <number>] [-INCLUDE <list>] [-EXCLUDE <list>] [-NODATE] [-NOTIME] [-NOPID] [-NOTID] [-NOSEQ] [-NOTYPE] [-NOID] [-HELP]
Options	<ul style="list-style-type: none">◆ PATH—Allows you to specify a path to the log file◆ path—Specify the full path to the log file◆ TYPE—Allows you to filter the messages that are displayed◆ number—Specify 1 to display warning and error messages or specify 2 to display warnings, errors, and notifications◆ INCLUDE—Only includes specified IDs. All other IDs will not be displayed in the output◆ EXCLUDE—Excludes specified IDs. Ignore the specified IDs and display all others◆ list—A comma-separated list of IDs or ID ranges that follows the INCLUDE and EXCLUDE switches. A space should separate the switch from the list but within the list, there should be no spaces. Ranges are specified with a begin and end number and separated with a dash (-).◆ NODATE—Does not display the date in the output◆ NOTIME—Does not display the time in the output◆ NOPID—Does not display the process ID in the output◆ NOTID—Does not display the thread ID in the output◆ NOSEQ—Does not display the sequence number in the output◆ NOTYPE—Does not display the message type number in the output◆ NOID—Does not display the LogViewer ID in the output◆ HELP—Displays the command options
Examples	<ul style="list-style-type: none">◆ LogViewer -type 2◆ LogViewer -include 200
Notes	The default setting is -type 2 which displays both type 1 and 2 messages.

Mirror Pause

Command	MIRROR PAUSE
Description	Pauses a mirror that is in progress
Syntax	<u>MIRROR</u> PAUSE <con_id *>
Options	<ul style="list-style-type: none">◆ con_id—Connection ID assigned to the source/target connection◆ *—Specifies all connection IDs.
Examples	<ul style="list-style-type: none">◆ mirror pause 1◆ mirror pause *

Mirror Resume

Command	MIRROR RESUME
Description	Resumes a paused mirror
Syntax	MIRROR RESUME <con_id *>
Options	<ul style="list-style-type: none">◆ con_id—Connection ID assigned to the source/target connection◆ *—Specifies all connection IDs.
Examples	<ul style="list-style-type: none">◆ <code>mirror resume 1</code>◆ <code>mirror resume *</code>

Mirror Start

Command	MIRROR START
Description	Initiates the mirror process
Syntax	MIRROR START <con_id> [CLEAN DIFFERENT [,NEWER] ,CHECKSUM NOCHECKSUM] [ORPHANS NOORPHANS] [CALCULATE NOCALCULATE] [CLEARRESTOREREQUIRED]
Options	<ul style="list-style-type: none">◆ con_id—Connection ID assigned to the source/target connection◆ CLEAN—Deletes all files in the specified target path prior to mirroring◆ DIFFERENT—Mirrors only those files that are different based on the file date, time, and/or size◆ NEWER—Mirrors only those files that are newer on the source than on the target◆ CHECKSUM—Mirrors only those blocks that are different based on block checksum comparisons◆ NOCHECKSUM—Does not perform a checksum comparison when mirroring files◆ ORPHANS—Removes orphan files on the target◆ NOORPHANS—Does not remove orphan files on the target◆ CALCULATE—Calculate the size of the replication set prior to mirroring◆ CLEARRESTOREREQUIRED—Clears the restore required flag and initiates the mirror
Examples	<ul style="list-style-type: none">◆ <code>mirror start 1 different, newer</code>◆ <code>mirror start 3 clean</code>◆ <code>mirror start 2 different, checksum</code>
Notes	<ul style="list-style-type: none">◆ In order to use the <code>clean</code> option, you must have the cleaner program setting enabled. For detailed information, see the cleaner option in the table Storage Mirroring Text Client program settings on page 14-15.◆ By default, orphan files will not be removed and the size of the replication set database will be calculated.

Mirror Stop

Command	MIRROR STOP
Description	Stops a paused mirror
Syntax	MIRROR STOP <con_id *>
Options	<ul style="list-style-type: none">◆ con_id—Connection ID assigned to the source/target connection◆ *—Specifies all connection IDs
Examples	<ul style="list-style-type: none">◆ mirror stop 1◆ mirror stop *

Monitor Account Add

Command	MONITOR ACCOUNT ADD
Description	Identifies a user account with update permissions within Active Directory
Syntax	MONITOR ACCOUNT ADD <username> <password> [<i>target_machine</i>]
Options	<ul style="list-style-type: none">◆ username—Name of the user in the format domain\user or user@domain◆ password—Password associated with username◆ target_machine—Name of the target machine
Examples	monitor account add administrator ***** jersey
Notes	<ul style="list-style-type: none">◆ The monitor account add command is not available when scrolling through the Text Client command history.◆ If characters in the password include non-alphanumeric characters, the password field must be enclosed in quotation marks.◆ The password cannot be a Storage Mirroring keyword. These are any DTCL commands (source, target, and so on) or any DTCL shortcut commands (env, mon, rep, and so on).◆ The password cannot be blank.

Monitor Account Remove

Command	MONITOR ACCOUNT REMOVE
Description	Removes the user account used to failover Active Directory hostnames
Syntax	MONITOR ACCOUNT REMOVE [<i>target_machine</i>]
Options	target_machine —Name of the target machine
Examples	monitor account remove jersey

Monitor Clear

Command	MONITOR CLEAR
Description	Clears all of the failover configuration and monitoring parameters for the specified monitor machine
Syntax	<u>MONITOR</u> CLEAR [<i>monitor</i>]
Options	<i>monitor</i> —Name of the source machine designated as the monitor
Examples	<code>monitor clear indy</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

Monitor Create

Command	MONITOR CREATE
Description	Establishes a source machine as a failover monitor. This is the machine that will be monitored by a target machine in case it should experience a failure.
Syntax	<u>MONITOR</u> CREATE < <i>source_machine</i> >
Options	<i>source_machine</i> —Name of the machine
Examples	<code>monitor create indy</code>

Monitor Delete

Command	MONITOR DELETE
Description	Deletes the specified failover monitor and all of its parameters
Syntax	<u>MONITOR</u> <u>DELETE</u> < <i>monitor</i> >
Options	<i>monitor</i> —Name of the source machine designated as the monitor
Examples	<code>monitor delete indy</code>
Notes	In order to successfully delete a monitor, the monitor must not be running on the server. Use the <code>monitor stop</code> command to ensure the monitor is not running.

Monitor Display

Command	<code>MONITOR DISPLAY</code>
Description	Displays the monitoring and failover configuration settings for the specified monitor machine
Syntax	<code><u>MONITOR</u> <u>DISPLAY</u> <<i>monitor</i>></code>
Options	<i>monitor</i> —Name of the source machine designated as the monitor
Examples	<code>monitor display indy</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

Monitor Get

Command	<code>MONITOR GET</code>
Description	Displays the machines that are currently being monitored by the specified target machine
Syntax	<code><u>MONITOR</u> <u>GET</u> <<i>target_machine</i>></code>
Options	<i>target_machine</i> —The name of the target machine
Examples	<code>monitor get jersey</code>
Notes	Because DTCL does not display interactive information, this command is only valid in the Text Client.

Monitor List

Command	<code>MONITOR LIST</code>
Description	Displays a list of all failover monitor machines
Syntax	<code><u>MONITOR</u> <u>LIST</u></code>

Monitor Move

Command `MONITOR MOVE`

Description Designates the IP address that will be failed over to the specified target NIC.

Syntax `MONITOR MOVE <IP_address> TO NIC <target_NIC>`
 `USE NONE`
 | `ICMP INTERVAL <interval> TIMEOUT <timeout>`
 | `SERVICE INTERVAL <interval> TIMEOUT <timeout>`
`[monitor] | <NOTEST> [monitor]`

Options

- ◆ **IP_address**—The IP address which should be moved during failover
- ◆ **target_NIC**—The integer value of the target NIC obtained from the `niclist` command
- ◆ **NONE**—Does not actively monitor the source machine. You will be responsible for identifying when a failure has occurred and initiating failover manually.
- ◆ **ICMP**—Monitors the source by using ICMP pings
- ◆ **INTERVAL interval**—The frequency, in seconds, of the monitor requests sent to the source machine to see if it is online and active
- ◆ **TIMEOUT timeout**—The number of seconds before failover will occur. This number is reset to its maximum each time the source sends a response to the monitor request.
- ◆ **SERVICE**—Monitors the source by using UDP heartbeats.
- ◆ **monitor**—Name of the source machine designated as the monitor
- ◆ **NOTEST**—Allows you to failover an IP address without sending monitor requests or expecting responses from the source. This option should only be used if you are monitoring multiple IP addresses but do not want to send monitor requests to each address.

Examples

- ◆ `monitor move 205.31.2.57 to nic 1 use icmp interval 5 timeout 25`
- ◆ `monitor move 205.31.2.68 to nic 2 use none notest`

Notes

- ◆ If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.
- ◆ If you choose to use UDP heartbeats (the service option), your source and target must be using the same Storage Mirroring port settings. If you choose to use ICMP pings (the ICMP option), the port settings can be different.

Monitor Option

Command MONITOR OPTION

Description Configures the settings to determine how failover will be performed

Syntax `MONITOR OPTION [MOVENAME | NOMOVENAME] [MOVEADDRESSES | NOMOVEADDRESSES] [MOVESHARES | NOMOVESHARES] [ADD | REPLACE] [, USESHAREFILE | NOUSESHAREFILE][, FAILONE | FAILALL] [, FODELAY | NOFODELAY] [, FBDELAY | NOFBDELAY] [, INTERVENTION | NOINTERVENTION] [, FOHOSTNAME | NOFOHOSTNAME] [, FBHOSTNAME | NOFBHOSTNAME][monitor]`

- Options**
- ◆ **MOVENAME**—Moves the server name during failover
 - ◆ **NOMOVENAME**—Does not move the server name during failover
 - ◆ **MOVEADDRESSES**—Moves the IP address(es) during failover
 - ◆ **NOMOVEADDRESSES**—Does not move the IP address(es) during failover
 - ◆ **MOVESHARES**—Moves the shares during failover
 - ◆ **NOMOVESHARES**—Does not move the shares during failover
 - ◆ **ADD**—Specifies the source machine's identity is added to the target machine's identity when failover occurs
 - ◆ **REPLACE**—Specifies that the source machine's identity replaces the target machine's identity when failover occurs
 - ◆ **USESHAREFILE**—Use the Storage Mirroring generated .shr file to determine shares
 - ◆ **NOUSESHAREFILE**—Do not use the Storage Mirroring generated .shr file to determine shares
 - ◆ **FAILONE**—When multiple IP addresses exist on a monitor machine, only the failed address is failed over to the target machine
 - ◆ **FAILALL**—When multiple IP addresses exist on a monitor machine, all of the addresses will fail over to the target machine even if only one address fails
 - ◆ **FODELAY**—Guarantees that the pre-failover script has completed before failing over
 - ◆ **NOFODELAY**—Does not guarantee that the pre-failover script has completed before failing over
 - ◆ **FBDELAY**—Guarantees that the pre-failback script has completed before failing back
 - ◆ **NOFBDELAY**—Does not guarantee that the pre-failback script has completed before failing back
 - ◆ **INTERVENTION**—Specifies that network administrator intervention is required before failover begins
 - ◆ **NOINTERVENTION**—Specifies that network administrator intervention is not required before failover begins
 - ◆ **FOHOSTNAME**—Removes the host SPN from Active Directory on the source and adds it to Active Directory on the target during failover
 - ◆ **NOFOHOSTNAME**—Does not modify Active Directory on the source or target during failover
 - ◆ **FBHOSTNAME**—Returns the host SPN on the source and target back to their original settings during failback
 - ◆ **NOFBHOSTNAME**—Does not modify Active Directory on the source or target during failback
 - ◆ ***monitor***—Name of the source machine designated as the monitor

Examples

- ◆ `monitor option`
- ◆ `monitor option nomovename`

Notes The default settings are movename, moveaddress, moveshares, add, usesharefile, failall, fodelay, fbdelay, intervention, nofohostname, and nofbhostname.

Monitor Remove

Command	MONITOR REMOVE
Description	Removes an IP address that is currently being monitored
Syntax	MONITOR REMOVE <IP_address> [<i>monitor</i>]
Options	<ul style="list-style-type: none">♦ <i>IP_address</i>—The currently monitored IP address that should be removed♦ <i>monitor</i>—Name of the source machine designated as the monitor
Examples	<ul style="list-style-type: none">♦ <code>monitor remove 205.31.2.57 indy</code>♦ <code>monitor remove 205.31.2.68</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

Monitor Script Add

Command	MONITOR SCRIPT ADD
Description	Specifies the scripts that should be run during the failover and failback processes
Syntax	MONITOR SCRIPT ADD <type> <script_name> [<i>monitor</i>]
Options	<ul style="list-style-type: none">♦ <i>type</i>—Any of the following script types:<ul style="list-style-type: none">♦ PREFAILOVER—Specifies that the file is a pre-failover script to be run on the target before failover occurs♦ POSTFAILOVER—Specifies that the file is a post-failover script to be run on the target after failover occurs♦ PREFAILBACK—Specifies that the file is a pre-failback script to be run on the target before failback occurs♦ POSTFAILBACK—Specifies that the file is a post-failback script to be run on the target after failback occurs♦ SRCPOSTFAILBACK—Specifies that the file is a post-failback script to be run on the source after failback occurs♦ <i>script_name</i>—Full path and name of the script file♦ <i>monitor</i>—Name of the source machine designated as the monitor
Examples	<code>monitor script add prefailback "C:\Program Files\OpenView\Storage Mirroring\preback.bat"</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

Monitor Script Remove

Command	<code>MONITOR SCRIPT REMOVE</code>
Description	Specifies the scripts that should not be run during the failover and failback processes
Syntax	<code><u>MONITOR</u> SCRIPT <u>REMOVE</u> <type> [<i>monitor</i>]</code>
Options	<ul style="list-style-type: none">♦ type—Any of the following script types:<ul style="list-style-type: none">♦ PREFAILOVER—Specifies that the file is a pre-failover script to be run on the target before failover occurs♦ POSTFAILOVER—Specifies that the file is a post-failover script to be run on the target after failover occurs♦ PREFAILBACK—Specifies that the file is a pre-failback script to be run on the target before failback occurs♦ POSTFAILBACK—Specifies that the file is a post-failback script to be run on the target after failback occurs♦ SRCPOSTFAILBACK—Specifies that the file is a post-failback script to be run on the source after failback occurs♦ monitor—Name of the source machine designated as the monitor
Examples	<ul style="list-style-type: none">♦ <code>monitor script remove prefailback</code>♦ <code>monitor script remove postfailover</code>
Notes	If you do not specify a monitor, the current source designated as the monitor will be used. If you have not identified a monitor, you will receive an error message stating that a monitor has not been selected.

Monitor Start

Command	<code>MONITOR START</code>
Description	The configured failover target starts monitoring the source machine that was designated as the monitor.
Syntax	<code><u>MONITOR</u> START <<i>monitor</i>> [ON <<i>target_machine</i>>]</code>
Options	<ul style="list-style-type: none">♦ monitor—Name of the source machine designated as the monitor♦ target_machine—Name of the target machine that will be monitoring the source machine
Examples	<code>monitor start indy on jersey</code>

Monitor Stop

Command	<code>MONITOR STOP</code>
Description	Stops monitoring the source machine for a failure
Syntax	<code><u>MONITOR</u> STOP <monitor> [ON <target_machine>]</code>
Options	<ul style="list-style-type: none">♦ monitor—Name of the source machine designated as the monitor♦ target_machine—Name of the target machine that will no longer be monitoring the source machine
Examples	<code>monitor stop indy on jersey</code>

Monitor Use

Command	<code>MONITOR USE</code>
Description	Specifies the source machine designated as the monitor that will be used in subsequent monitor commands
Syntax	<code><u>MONITOR</u> USE <monitor></code>
Options	monitor —Name of the source machine designated as the monitor
Examples	<code>monitor use indy</code>

Niclist

Command	<code>NICLIST</code>
Description	Displays the NICs available on the specified target machine. Each NIC is assigned an integer value and this value is used in the <code>monitor move</code> command.
Syntax	<code>NICLIST [<i>target_machine</i>]</code>
Options	target_machine —Name of the machine
Examples	<code>niclist jersey</code>
Notes	<ul style="list-style-type: none">♦ If you do not specify a machine name, the value from the current target will be returned. If you have not identified a target, no data will be returned.♦ If you have not logged into the target machine, no data will be displayed.

Orphans Disable

Command	<code>ORPHANS DISABLE</code>
Description	Disables orphan file removal during subsequent mirror or verification processes.
Syntax	<code>ORPHANS DISABLE <i>con_id</i></code>
Options	<i>con_id</i> —Connection ID assigned to the source/target connection
Examples	<code>orphans disable 1</code>

Orphans Enable

Command	<code>ORPHANS ENABLE</code>
Description	Enables orphan file removal during subsequent mirror or verification processes.
Syntax	<code>ORPHANS ENABLE <i>con_id</i></code>
Options	<i>con_id</i> —Connection ID assigned to the source/target connection
Examples	<code>orphans enable 1</code>

Orphans Start

Command	<code>ORPHANS START</code>
Description	Manual operation to remove any orphan files on the target
Syntax	<code>ORPHANS START <i>con_id</i> [CLEARRESTOREREQUIRED]</code>
Options	<ul style="list-style-type: none">♦ <i>con_id</i>—Connection ID assigned to the source/target connection♦ <code>CLEARRESTOREREQUIRED</code>—Clears the restore required flag and initiates the orphan operation
Examples	<code>orphans start 1</code>

Orphans Stop

Command	<code>ORPHANS STOP</code>
Description	Stops the process to remove orphan files on the target
Syntax	<code>ORPHANS STOP <i>con_id</i></code>
Options	<i>con_id</i> —Connection ID assigned to the source/target connection
Examples	<code>orphans stop 1</code>

Pause Target

Command	PAUSETARGET
Description	Allows you to pause the execution of Storage Mirroring operations on the target
Syntax	PAUSETARGET <target_machine> [FROM <source_machine>]
Options	<ul style="list-style-type: none">♦ target_machine—The name of the target machine where you want to pause execution of the Storage Mirroring operations♦ source_machine—The name of the source machine that is connected to the target
Examples	pausetarget jersey
Notes	<ul style="list-style-type: none">♦ You must be logged on to the target machine for this command to work.♦ If the target machine has not been identified using the <code>target</code> command, you must specify the <target_name> in the <code>pausetarget</code> command.

Ping

Command	PING
Description	Checks a specified machine to see if Storage Mirroring is running
Syntax	PING <machine>
Options	machine —Name of the machine
Examples	ping indy

QueueTask

Command	QUEUETASK
Description	Queues tasks inline with replication data.
Syntax	<u>QUEUETASK</u> <job_name> TO <target_machine> <u>ONQUEUE</u> = <task> [args] <u>ONTRANSMIT</u> = <task> [args] <u>ONRECEIVE</u> = <task> [args] <u>ONEXECUTE</u> = <task> [args] [<u>TIMEOUT</u> = <timeout>] [INTERACT NOINTERACT]

Options

- ◆ **job_name**—Any name that you have assigned to the job. This will be the identifier you will see in the log files.
- ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name. The target is required even if you are only queuing a task to be executed on the source.
- ◆ **ONQUEUE**—Execute the specified task on the source machine as soon as the source receives and queues the task. During heavy replication, there may be a delay while the task is queued inline with the replication operations.
- ◆ **ONTRANSMIT**—Execute the specified task on the source machine just before the source transmits the task to the target.
- ◆ **ONRECEIVE**—Execute the specified task on the target machine as soon as the target receives and queues the task.
- ◆ **ONEXECUTE**—Execute the specified task on the target when the target processes the task from the queue. Since the task is not executed until it is processed, if the target is paused, the task will be held in queue.
- ◆ **task**—The path and filename of the task to run relative to the machine it will be run on. Tasks include any valid executable or batch file. If the path contains non-alphanumeric characters, you must enclose it in quotation marks. The executables or files must exist in the specified location on the machine where they will be executed
- ◆ **args**—Arguments or options which need to be supplied with the task. Multiple arguments can be supplied in a space-separated list enclosed in quotation marks.
- ◆ **TIMEOUT**—Length of time to pause while waiting for the task to complete
- ◆ **timeout**—Any valid number followed by an optional time indicator indicating length of time. The valid time indicators include seconds, minutes, hours, and days. If you do not specify a time indicator, it will default to seconds. The number zero (0) indicates there is no timeout delay and the next operation is immediately processed. The keyword **FOREVER** indicates that the next operation is not processed until the task has completed execution. No **TIMEOUT <timeout>** setting will default to forever.
- ◆ **INTERACT**—Tasks interact with the desktop and, therefore, display on screen and run in the foreground
- ◆ **NOINTERACT**—Tasks do not interact with the desktop

Examples

```
qtask backup to jersey onrec=PauseAndBackup.bat onexec=resume.bat
```

Notes

- ◆ The default setting for this command is **nointeract**.
- ◆ Any combination of one or more execution points can be used with the same **queuetask** command.
- ◆ All processing messages, including errors, can be viewed in the Storage Mirroring logs and the Windows Event Viewer.
- ◆ **Onqueue** will still execute as soon as the task is placed on the queue even if transmission is stopped (manually stopped or paused, unmet scheduled transmission criteria, etc.). Any other option will not execute until transmission is restarted.
- ◆ If your source is in a restore required state, any task placed on the queue will be executed immediately. Use caution when submitting tasks while in this state so that the target does not get inadvertently updated.
- ◆ A task may be discarded in the following scenarios: all connections to a target are manually disconnected, replication is stopped for all connections to a target, or an auto-disconnect occurs.
- ◆ If a task is submitted after replication is stopped, the task will be executed immediately.
- ◆ If you disable task command processing while tasks are in queue, those tasks will not be executed.
- ◆ The user submitting the task command must be a member of the **Double-Take Admin** security group on both the source and target and the Storage Mirroring service must have proper privileges to access the files or run the commands specified in the task.

Quit

Command	QUIT
Description	Quits the Text Client and the Command Line Interactive client
Syntax	QUIT

Replication Start

Command	REPLICATION START
Description	Initiates the replication process
Syntax	<u>REPLICATION</u> START <conid *> [CLEARRESTOREREQUIRED]
Options	<ul style="list-style-type: none">◆ conid—Connection ID assigned to the source/target connection◆ *—Specifies all connection IDs◆ CLEARRESTOREREQUIRED—Clears the restore required flag and initiates replication
Examples	<ul style="list-style-type: none">◆ replication start 1◆ replication start *

Replication Stop

Command	REPLICATION STOP
Description	Stops the replication process
Syntax	<u>REPLICATION</u> STOP <conid *>
Options	<ul style="list-style-type: none">◆ conid—Connection ID assigned to the source/target connection◆ *—Specifies all connection IDs
Examples	<ul style="list-style-type: none">◆ replication stop 1◆ replication stop *
Notes	<p>If you stop replication manually, you will need to perform a remirror when it is restarted.</p> <p>If you only need to temporarily pause data from being transmitted to the target and do not want to stop replication, use the <code>transmission pause</code> command. When you are ready to restart the transmission, use the <code>transmission resume</code> command.</p>

Repset Calculate

Command	REPSET CALCULATE
Description	Calculates the size of a replication set
Syntax	REPSET <u>CALCULATE</u> [<i>repset</i>]
Options	<i>repset</i> —Name of the replication set
Examples	<ul style="list-style-type: none">♦ repset calculate marketing♦ repset calc "Exchange repset"
Notes	<ul style="list-style-type: none">♦ If a replication set name is not specified, the active replication set will be used.♦ If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.♦ The results of the calculation are logged to the Storage Mirroring log file.

Repset Create

Command	REPSET CREATE
Description	Creates a replication set
Syntax	REPSET CREATE < <i>name</i> >
Options	<i>name</i> —Name of the replication set
Examples	<ul style="list-style-type: none">♦ repset create "Exchange Repset"♦ repset create marketing
Notes	<ul style="list-style-type: none">♦ If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.♦ If the name of the replication set is the same as a word used in a DTCL command (for example, repset create repset), you will receive an error. If you want to use a word like this, you must enclose it in quotation marks (repset create "repset").

Repset Delete

Command	REPSET DELETE
Description	Deletes the specified replication set
Syntax	REPSET <u>DELETE</u> < <i>repset</i> >
Options	<i>repset</i> —Name of the replication set
Examples	repset delete "Exchange Repset"
Notes	If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.

Repset Display

Command	<code>REPSET DISPLAY</code>
Description	Displays the rule of a replication set
Syntax	<code>REPSET <u>DISPLAY</u> [<i>repset</i>]</code>
Options	<i>repset</i> —Name of the replication set
Examples	<ul style="list-style-type: none">♦ <code>repset display</code>♦ <code>repset display Exchange</code>
Notes	<ul style="list-style-type: none">♦ If you do not specify a replication set name, the current replication set will be used.♦ If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.

Repset List

Command	<code>REPSET LIST</code>
Description	Lists all replication set names for the currently selected source
Syntax	<code>REPSET LIST</code>

Repset Resync

Command	<code>REPSET RESYNC</code>
Description	Retrieves the last saved replication set settings from the source.
Syntax	<code>REPSET RESYNC</code>

Repset Rule Add

Command	REPSET RULE ADD
Description	Adds a rule to a replication set. A rule is the specification of a path including volume, directories, wild cards, and/or file names.
Syntax	REPSET RULE ADD <path> [<u>I</u> NCLUDE <u>E</u> XCLUDE] [, <u>R</u> ECURSIVE <u>N</u> ONRECURSIVE] [TO <repset>]
Options	<ul style="list-style-type: none">◆ path—Volume, directory, wild card, and/or file name◆ INCLUDE—Include the specified path in the replication set◆ EXCLUDE—Exclude the specified path in the replication set◆ RECURSIVE—All subdirectories and files are recursively included or excluded◆ NONRECURSIVE—All subdirectories and files are non-recursively included or excluded◆ repset—Name of the replication set
Examples	repset rule add c:\exchange to "Exchange Repset"
Notes	<ul style="list-style-type: none">◆ The default settings for this command are include and recursive.◆ If you do not specify a replication set name, the current replication set will be used.◆ If the path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.◆ Verify what files can be included by reviewing Multi-Client Access on page 5-5.

Repset Rule Remove

Command	REPSET RULE REMOVE
Description	Removes a rule from a replication set
Syntax	REPSET RULE <u>R</u> EMOVE <path> [<u>F</u> ROM <repset>]
Options	<ul style="list-style-type: none">◆ path—Volume, directory, wild card, and/or file name◆ repset—Name of the replication set
Examples	repset rule remove c:\exchange from "Exchange Repset"
Notes	<ul style="list-style-type: none">◆ If you do not specify a replication set name, the current replication set will be used.◆ If the path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.

Repset Save

Command	<code>REPSET SAVE</code>
Description	Saves all replication set rules for the currently selected source
Syntax	<code>REPSET SAVE</code>

Repset Use

Command	<code>REPSET USE</code>
Description	Specifies a replication set as the active replication set
Syntax	<code>REPSET USE <repset></code>
Options	<i>repset</i> —Name of the replication set
Examples	<ul style="list-style-type: none">◆ <code>repset use "Exchange Repset"</code>◆ <code>repset use marketing_database</code>
Notes	If the replication set name includes non-alphanumeric characters, the name must be enclosed in quotation marks.

Restore

Command	RESTORE
Description	Initiates the restoration process
Syntax	<pre>RESTORE <repset> FROM <target_machine> [ORIGINAL <original_source_machine>] [, OVERWRITE NOOVERWRITE][, OVERWRITENEWER NOOVERWRITENEWER] [, USETARGETDB NOUSETARGETDB][, RESTOREDBTOO NORESTOREDBTOO] [, CHECKSUM NOCHECKSUM][, ORPHANS NOORPHANS]</pre>
Options	<ul style="list-style-type: none">◆ repset—Name of the replication set◆ target_machine—Name of the target machine◆ original_source_machine—Specifies that the replication set was not originated on the current source machine◆ OVERWRITE—Overwrites files on the source◆ NOOVERWRITE—Does not overwrite files on the source◆ OVERWRITENEWER—Overwrites files on the source even if the source file is newer than on the target◆ NOOVERWRITENEWER—Does not overwrite files on the source that are newer on the source than on the target◆ USETARGETDB—Uses the replication set from the target machine◆ NOUSETARGETDB—Uses the replication set from the source machine◆ RESTOREDBTOO—Restores the replication set from the target to the source◆ NORESTOREDBTOO—Does not restore the replication set from the target to the source◆ CHECKSUM—Performs a block checksum comparison and only restores those blocks that are different◆ NOCHECKSUM—Does not perform a block checksum comparison and restores those files that are different◆ ORPHANS—Removes orphan files on the target◆ NOORPHANS—Does not remove orphan files on the target
Examples	<pre>restore "Exchange Repset" from jersey, overwritenewer, usetargetdb</pre>
Notes	<ul style="list-style-type: none">◆ The default settings for this command are <code>overwrite</code>, <code>overwritenewer</code>, <code>usetargetdb</code>, <code>restoredbtoo</code>, and <code>noorphans</code>.◆ The source command is required before each use of the restore command.◆ If a path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.◆ If you are using the <code>restore</code> command with the DTCL interactive or batch client (<code>dtcl -i</code> or <code>dtcl -f</code>), the command requires the original source option. This option is not required if you are using the Text Client.

Resume Target

Command	<code>RESUMETARGET</code>
Description	Allows you to resume the execution of Storage Mirroring operations on the target
Syntax	<code>RESUMETARGET <target_machine> [FROM <source_machine>]</code>
Options	<ul style="list-style-type: none">♦ target_machine—The name of the target machine where you want to resume execution of the Storage Mirroring operations♦ source_machine—The name of the source machine from where you want to resume execution of the Storage Mirroring operations
Examples	<code>resumetarget jersey</code>
Notes	<ul style="list-style-type: none">♦ You must be logged on to the target machine for this command to work.♦ If the target machine has not been identified using the <code>target</code> command, you must specify the <code><target_machine></code> in the <code>resumetarget</code> command.

Schedule Clear

Command	<code>SCHEDULE CLEAR</code>
Description	Clears the existing transmission schedule for the specified or current target machine
Syntax	<code>SCHEDULE <target_machine> CLEAR</code>
Options	target_machine —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>schedule jersey clear</code>

Schedule Disable

Command	<code>SCHEDULE DISABLE</code>
Description	Disables the transmission schedule without clearing the schedule data
Syntax	<code>SCHEDULE <target_machine> DISABLE</code>
Options	target_machine —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>schedule jersey disable</code>

Schedule Enable

Command	<code>SCHEDULE ENABLE</code>
Description	Enables the transmission schedule
Syntax	<code><u>SCHEDULE</u> <target_machine> ENABLE</code>
Options	target_machine —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>schedule jersey enable</code>

Schedule End

Command	<code>SCHEDULE END</code>
Description	Sets criteria to end the transmission of data from the source to the target
Syntax	<code><u>SCHEDULE</u> <target_machine> END [DURATION = <number> <time_units>][BYTES = <bytes>]</code>
Options	<ul style="list-style-type: none">♦ target_machine—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.♦ number—Any number indicating the length of time before the transmission ends♦ time_units—Minutes (min), hours (hr), or days (day)♦ bytes—Number of bytes transmitted before the transmission ends
Examples	<ul style="list-style-type: none">♦ <code>schedule jersey end duration=3 hr</code>♦ <code>schedule philly end bytes=1500000</code>
Notes	If you use both of the end options, <code>duration</code> and <code>bytes</code> , the transmission will be stopped when the first end option value is met.

Schedule Start

Command	<code>SCHEDULE START</code>
Description	Sets criteria to start the transmission of data from the source to the target
Syntax	<code>SCHEDULE <target_machine> START [STARTTIME = <mm/dd/yy> <hh:mm>] [MEMLIMIT = <percent>] [QUEUESIZE = <bytes>] [EVERY <number> <time_units>]</code>
Options	<ul style="list-style-type: none">◆ target_machine—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.◆ mm/dd/yy—Date in month/day/year format indicating when the transmission will begin◆ hh:mm—Time in hour:minute format using the 24 hour clock indicating when the transmission will begin◆ percent—Any number between 0 and 100 indicating the percentage of the disk queue and system memory that must be in use to initiate the transmission process◆ bytes—Number of bytes that must be in the source queue and system memory to initiate the transmission process◆ number—Any number indicating how often the transmission process will be repeated◆ time_units—Minutes (min), hours (hr), or days (day)
Examples	<ul style="list-style-type: none">◆ <code>schedule jersey start starttime=3/11/04 03:30</code>◆ <code>schedule philly start queuesize=400000</code>◆ <code>schedule jersey start starttime=3/03/04 18:00, every 8 hr</code>◆ <code>schedule philly start memlimit=25, queuesize=150000000</code>
Notes	<ul style="list-style-type: none">◆ The start option EVERY cannot be used by itself and cannot be the first option in a string of options.◆ If you use more than one start option, the transmission will begin when the first start option value is met. Additionally, each option after the first must be separated by a comma, as illustrated in the Examples.

Schedule Window

Command	<code>SCHEDULE WINDOW</code>
Description	Sets criteria to only allow transmissions during a certain period of time
Syntax	<code>SCHEDULE <target_machine> WINDOW <hh:mm> TO <hh:mm></code>
Options	<ul style="list-style-type: none">◆ target_machine—Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.◆ hh:mm—Time in hour:minute format using the 24 hour clock. The first time is when the transmission will begin and the second time is when the transmission will end.
Examples	<code>schedule jersey window 23:00 to 06:00</code>
Notes	Establishing a transmission window by itself is not sufficient to start a transmission. You will need to specify a start criteria.

Set

Command	SET
Description	The Text Client/DTCL sends a request to Storage Mirroring to modify the value of a Storage Mirroring program setting. The setting is immediately applied and is then forwarded to the registry where it is also applied.
Syntax	SET <option>=<value> [<i>machine</i>]
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>machine</i>—Name of the machine
Examples	<ul style="list-style-type: none">♦ set netport=1100♦ set VerifyLogName="HA server.log"♦ set DefaultAddress="129.12.34.3"
Notes	<ul style="list-style-type: none">♦ Some settings, although immediately applied to Storage Mirroring, will not take effect until the client or service is restarted. For example, the Text Client must be closed and reopened for a change to ClientLogName to take effect. And the service must be restarted to cause a change in the modules loaded if the LoadSourceTarget setting is changed.♦ If you do not specify a machine name, the value from the current source will be updated. If you have not identified an active source, no changes will be made.

Set Local

Command	SETLOCAL
Description	The Text Client/DTCL sends a request directly to the registry, bypassing Storage Mirroring, to modify the value of a Storage Mirroring program setting.
Syntax	SETLOCAL <option>=<value>
Options	<ul style="list-style-type: none">♦ <i>option</i>—See Storage Mirroring Text Client program settings on page 14-15.♦ <i>value</i>—See Storage Mirroring Text Client program settings on page 14-15.
Examples	<ul style="list-style-type: none">♦ setlocal netport=1100♦ setlocal VerifyLogName="HA server.log"♦ setlocal DefaultAddress="129.12.34.3"
Notes	<ul style="list-style-type: none">♦ Setlocal is designed to allow you to edit DTCL and Text Client program settings on a machine where the Storage Mirroring service is not running. Therefore, if you are running the Storage Mirroring service, any changes made with setlocal will not take effect until the service is restarted.

Shutdown

Command	SHUTDOWN
Description	Stops the Storage Mirroring service
Syntax	SHUTDOWN < <i>machine</i> >
Options	<i>machine</i> —Name of the machine
Examples	shutdown indy

Source

Command	SOURCE
Description	Identifies a machine as the active source machine
Syntax	<u>SOURCE</u> < <i>source_machine</i> >
Options	<i>source_machine</i> —Name of the machine
Examples	source indy

Statslog Start

Command	STATSLOG START
Description	Starts the DTStats statistics logging utility.
Syntax	STATSLOG START < <i>machine</i> > [TO < <i>filename</i> >] [EVERY < <i>minute</i> >] [MAXSIZE < <i>kilobytes</i> >]
Options	<ul style="list-style-type: none">♦ <i>machine</i>—Name of the machine♦ <i>filename</i>—Any valid path and filename to which the statistical information will be logged♦ <i>minutes</i>—The number of minutes to write to the log file♦ <i>kilobytes</i>—The maximum file size in kilobytes
Examples	statslog start indy to dtstat.sts every 20 maxsize 200000
Notes	<ul style="list-style-type: none">♦ The default settings are <code>statistic.sts</code>.♦ The default interval is every 5 minutes.♦ The default file size is 10485760 bytes (10 MB).

Statslog Status

Command	<code>STATSLOG STATUS</code>
Description	Identifies whether or not the DTStats statistics logging utility is currently running
Syntax	<code>STATSLOG STATUS <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>statslog status indy</code>

Statslog Stop

Command	<code>STATSLOG STOP</code>
Description	Stops the DTStats statistics logging utility.
Syntax	<code>STATSLOG STOP <machine></code>
Options	<i>machine</i> —Name of the machine
Examples	<code>statslog stop indy</code>

Status

Command	<code>STATUS</code>
Description	Requests connection and statistical information
Syntax	<code>STATUS <u>C</u>ONNECT MIRROR <u>R</u>EPLICATE <u>T</u>RANSMIT <con_id></code>
Options	<ul style="list-style-type: none">◆ <code>CONNECT</code>—Displays connection information◆ <code>MIRROR</code>—Displays mirroring information◆ <code>REPLICATE</code>—Displays replication information◆ <code>TRANSMIT</code>—Displays transmission information◆ <i>con_id</i>—Connection ID assigned to the source/target connection
Examples	<ul style="list-style-type: none">◆ <code>status connect 1</code>◆ <code>status replicate 1</code>

Target

Command	TARGET
Description	Identifies a machine as the active target machine
Syntax	<u>TARGET</u> <target_machine>
Options	<i>target_machine</i> —Name of the machine
Examples	target jersey
Notes	You must be logged into a machine using the login command before using the target command.

Testcon

Command	TESTCONNECTIONS
Description	Sends a test command to determine if Storage Mirroring is running. If there are active connections, the replication set(s), IP address(es), and connection ID(s) will be included in the testcon log file.
Syntax	<u>TESTCONNECTIONS</u> <machine> [<i>filename</i> [OVERWRITE]]
Options	<ul style="list-style-type: none">♦ <i>machine</i>—The name of the machine♦ <i>filename</i>—The name of the log file♦ OVERWRITE—Indicates that existing data in the file will be overwritten
Examples	testcon indy connstat
Notes	<ul style="list-style-type: none">♦ The default log file is status.dts.♦ The log file is located in the same directory as the other Storage Mirroring log files.♦ The file extension .dts is appended to any file name supplied.♦ By default, the log file is not overwritten.

TimeNow

Command	TIMENOW
Description	Prints the current date and time.
Syntax	<u>TIMENOW</u>
Examples	timenow
Notes	This command is useful for scripting to determine process start and stop times.

Transmission Pause

Command	TRANSMISSION PAUSE
Description	Pauses the transmission pause
Syntax	<u>TRANSMISSION PAUSE</u> < <i>target_machine</i> >
Options	<i>target_machine</i> —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>transmission pause jersey</code>

Transmission Resume

Command	TRANSMISSION RESUME
Description	Resumes a paused transmission
Syntax	<u>TRANSMISSION RESUME</u> < <i>target_machine</i> >
Options	<i>target_machine</i> —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>transmission resume jersey</code>

Transmission Start

Command	TRANSMISSION START
Description	Initiates the transmission process
Syntax	<u>TRANSMISSION START</u> < <i>target_machine</i> >
Options	<i>target_machine</i> —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>transmission start jersey</code>

Transmission Stop

Command	<code>TRANSMISSION STOP</code>
Description	Stops the transmission process
Syntax	<code><u>TRANSMISSION STOP</u> <target_machine></code>
Options	target_machine —Name of the target machine, an IP address on the target machine, or a virtual IP address. Depending on the network route Storage Mirroring uses to access the target, you may need to specify the fully qualified domain name.
Examples	<code>transmission stop jersey</code>

Unload Source

Command	<code>UNLOAD SOURCE</code>
Description	Unloads the Storage Mirroring source module
Syntax	<code>UNLOAD SOURCE <machine></code>
Options	machine —Name of the machine
Examples	<code>unload source indy</code>

Unload Target

Command	<code>UNLOAD TARGET</code>
Description	Unloads the Storage Mirroring target module
Syntax	<code>UNLOAD TARGET <machine></code>
Options	machine —Name of the machine
Examples	<code>unload target jersey</code>

Verify

Command	VERIFY
Description	Verifies the integrity of the data between the source and target machines.
Syntax	VERIFY <con_id> [SYNC [,NEWER][,CHECKSUM NOCHECKSUM] NOSYNC] [,CHECKSUM NOCHECKSUM] [STARTTIME=<mm/dd/yy> [hh:mm]] [EVERY <number> <time_units>] [ORPHANS NOORPHANS] [CLEARRESTOREREQUIRED]
Options	<ul style="list-style-type: none">◆ con_id—The connection ID which should be verified◆ SYNC—Synchronizes any data that is different◆ NEWER—Mirrors only those files that are newer on the source than on the target◆ CHECKSUM—Mirrors those blocks that are different on the source than on the target based on checksum comparisons◆ NOCHECKSUM—Does not perform a checksum comparison when mirroring files◆ NOSYNC—Do not synchronize any data that is different◆ STARTTIME—Starts the verification process at the time specified◆ mm/dd/yy—Date in month/day/year format when the verification process will begin◆ hh:mm—Time in hour:minute format using the 24 hour clock when the verification process will begin◆ EVERY—Repeat the verification process at the frequency specified◆ number—Length of time to repeat the verification process◆ time_units—Minutes (min), hours (hr), or days (day)◆ ORPHANS—Removes orphan files on the target◆ NOORPHANS—Does not remove orphan files on the target◆ CLEARRESTOREREQUIRED—Clears the restore required flag and initiates the verification
Examples	<ul style="list-style-type: none">◆ verify 1◆ verify 2 sync, newer◆ verify 2 every 2 hr
Notes	The default verification settings are sync, checksum, and noorphans.

Version

Command	VERSION
Description	Displays the Storage Mirroring version currently installed
Syntax	VERSION [MACHINE]
Options	MACHINE —Name of the machine
Examples	version indy
Notes	If no machine name is specified, the version of the active source will be displayed. If there is no active source, no version information will be displayed.

Wait

Command	<code>WAIT</code>
Description	This command is used in scripts to force the script to stop executing until the specified number of milliseconds has elapsed.
Syntax	<code>WAIT <ms></code>
Options	<code>ms</code> —Length of time in milliseconds
Examples	<code>wait 5000</code>
Notes	This sample waits for 5 seconds.

Wait on Mirror

Command	<code>WAITONMIRROR</code>
Description	This command is used in scripts to force the script to stop executing until the connection has finished mirroring or verifying.
Syntax	<code>WAITONMIRROR <cond_id></code>
Options	<code>con_id</code> —Connection ID assigned to the source/target connection
Examples	<code>waitonmirror 1</code>

Wait on Restore

Command	<code>WAITONRESTORE</code>
Description	This command is used in scripts to force the script to stop executing until the connection has finished restoring.
Syntax	<code>WAITONRESTORE <rest_id></code>
Options	<code>rest_id</code> —Restoration ID assigned to the source/target connection
Examples	<code>waitonrestore 1</code>
Notes	This command requires using the <code>source</code> command immediately before the <code>waitonrestore</code> command.

Wait on Target

Command	WAITONTARGET
Description	This command is used in scripts to force the script to stop executing until the target queue has been empty for the time specified.
Syntax	WAITONTARGET <target> [FROM <source>] <time>
Options	<ul style="list-style-type: none">♦ target—The name or IP address of the target to wait on♦ source—The name or IP address of the source♦ time—The number of seconds that the target queue needs to be empty before the command will return
Examples	<pre>waitontarget jersey 30</pre>

Write

Command	WRITE
Description	Displays the value of a DTCL variable
Syntax	WRITE \$<variable_name>
Options	variable_name —The name of the variable that you have established and want to display its value
Examples	<ul style="list-style-type: none">♦ <pre>write \$thetarget</pre>♦ <pre>write \$dbrepset</pre>

DTCL Scripting

To simplify using Storage Mirroring and DTCL, you can create script files to execute series and combinations of commands to meet specific needs.

NOTE: When creating DTCL scripts, each commented line must start and end with the pound (or number) sign (#). Each DTCL command and assignment must end with a semi-colon. However, no semicolon is necessary after the END keyword for the IF conditional and the loop statements.

When using the Command Line Client, a script file can be executed using the `-f` option. see [The Command Line Client](#) on page 3-3.

Variables

Variables can be used in commands to replace items that vary such as machine or domain names. A variable name must start with the `$` character and can contain letters, numbers, strings, and integers. Values are assigned to variables using the '=' assignment statement. For example, valid variable assignment statements are:

```
$MyDomain="domain.com";  
$MySource=indy;
```

All DTCL commands return values which can also be assigned to variables. For example, the `connection` command returns a connection ID for the new connection. The statement:

```
$ConnectionID=CONNECT dataset to target;
```

saves the connection ID as a variable. The variable components of the connect command could be replaced as well. For example:

```
$TheRepset="Exchange rep set";  
$TheTarget=jersey;  
$ConnectionID=CONNECT $TheRepset to $TheTarget;
```

Flow control

DTCL scripts can take advantage of four programming constructs.

- ◆ **IF conditional**—The IF conditional expression is a comparison between two values or variables. Use the following syntax for an IF conditional:

```
IF <relational expression> THEN <statement_block>
[ ELSE <statement_block> ] END
```

Use the following conditions to create the relational expression:

- ◆ **=** Is equal to
 - ◆ **!=** Is not equal to
 - ◆ **<** Is less than
 - ◆ **<=** Is less than or equal to
 - ◆ **>** Is greater than
 - ◆ **>=** Is greater than or equal to
 - ◆ **statement block** Any sequence of valid DTCL commands
- ◆ **FOR Loop**—The FOR loop sets a variable to a start value, executes each statement in the statement block, and then adds the step value to the variable. If the new value of the variable does not exceed the end value then the statements will be executed again. This continues until the variable exceeds the end value. If a step value is not provided, the default adds one to the variable. Use the following syntax for the FOR loop:

```
FOR <variable> = <start_value> TO <end_value>
[STEP <step_value>] DO <statement_block> END
```
 - ◆ **WHILE Loop**—The WHILE loop evaluates a relational expression and, if it is true, then the statement block is executed. When the statement block has completed execution, the expression will be reevaluated again and, if it is true, the statement block is executed again. This continues until the expression is false. Use the conditions specified in the IF conditional to create the relational expression and the following syntax for the WHILE loop:

```
WHILE <relational_expression> DO <statement_block> END
```
 - ◆ **WRITE Command**—The WRITE command writes values to the screen. These values can be variables or constants of type integer\$, string\$, date\$, or time\$. The write command is used in conjunction with either the IF conditional, WHILE loop, or FOR loop.

DTCL Return Values

Understanding the values returned by the DTCL commands will make variables and script flow control more useful.

NOTE: If the `connect` command returns a positive number, that number is the connection ID assigned to that connection.

The `login` command may return one of three positive values: 0 (no access granted), 1 (monitor access granted), or 2 (full access granted).

The following table includes all of the DTCL return values and a short description of the cause of the error.

NOTE: Some of the DTCL return values are internally used between servers and will not be displayed through the DTCL client.

Error Code	Description
-1	An unknown error code (generated when a command failed but the failure is not linked to a pre-defined error code)
-101	An invalid parameter was supplied in the DTCL command.
-102	The command entered is not a valid DTCL command or the syntax is incorrect.
-103	The machine being identified as the source does not have the Storage Mirroring source module loaded.
-104	No machine has been identified as a Storage Mirroring source.
-105	The machine being identified as the target does not have the Storage Mirroring target module loaded.
-106	The connection you are attempting to create has already been established.
-107	The connection you are attempting to manipulate does not exist.
-108	The mirror you are attempting to stop is not currently active.
-109	The machine identified does not exist or could not be located.
-110	The machine you are attempting to access is not responding.
-111	Storage Mirroring is running on the machine specified.
-112	An unknown connection error has occurred.
-113	The mirror you are attempting to start or resume is already active.
-114	The date specified is invalid. The valid format is mm/dd/yy.
-115	The time specified is invalid. The valid format is hh:mm.
-116	An invalid option was supplied in the DTCL command.
-117	The mirror you are attempting to resume is not paused.
-118	The connection you are attempting to resume is not paused.
-119	The connection you are attempting to manipulate does not exist.
-120	The connection you are attempting to establish is already connected.
-121	The mirror you are attempting to stop or pause is not currently running.
-122	The replication set you are attempting to create already exists.
-123	The replication set you are attempting to delete does not exist.
-124	The replication set you are attempting to manipulate does not exist.
-125	The connection for which you are attempting to start or resume replication is already replicating.

Error Code	Description
-126	The connection for which you are attempting to stop or pause replication is not replicating.
-127	The replication set you are attempting to enable is already enabled.
-128	The schedule you are attempting to initiate has not been defined.
-129	Multiple Management Console instances have modified the same replication set.
-130	The replication set you are attempting to change is already active.
-131	No machine has been identified as a Storage Mirroring target.
-132	Memory on the machine is low.
-133	Memory on the machine is sufficient.
-134	Replication is pending.
-135	An invalid option was supplied in the DTCL command.
-136	The replication set rule you are attempting to modify or delete does not exist.
-137	The mirror queue is full.
-138	The DTCL command used is not allowed because of insufficient security access.
-139	The DTCL schedule command used is invalid.
-140	The source path specified is invalid.
-141	You are attempting to save a replication set to which there have been no changes.
-142	The source machine specified cannot be changed because of insufficient security access.
-143	An invalid DTStats log file has been specified.
-144	The replication set you are attempting to save was not successfully saved.
-145	The connection you are attempting to establish failed because of insufficient security access.
-146	The cleaner option used with the mirror start clean command is disabled.
-147	The target mirror capacity high threshold has been met.
-148	The target mirror capacity low threshold has been met.
-149	The text client may need to be restarted for the new option to take effect.
-150	The target has been restarted.
-151	Replication is active. The machine is out of memory.
-152	Write access is blocked on the volume.
-153	The compression level specified is not supported or the source or target you are attempting to use do not support compression.
-153	The transmission you are attempting to pause is already paused.
-154	The transmission you are attempting to resume is already active.
-155	The DTCL command is not supported on the target specified.
-156	The DTCL command conversion for a varying Storage Mirroring version has failed.
-157	The source and target servers have incompatible Storage Mirroring versions.
-158	The source and target servers have incompatible operating system versions.
-159	NAS server to non-NAS server is an unsupported configuration.
-160	The Storage Mirroring target module is not loaded on the specified server.

Error Code	Description
-162	The target you are attempting to pause is already paused.
-163	The target you are attempting to pause is already attempting to pause but is pending.
-164	The target you are attempting to resume is already active.
-165	The target is retrying operations.
-166	The target is no longer retrying operations.
-167	Communication between the source and target is unavailable, therefore, Storage Mirroring could not determine if a restore is required.
-168	The source you are attempting to specify for failover is a replication only source and cannot be used for failover.
-169	The specified source could not be logged into during failover configuration.
-170	The task command processing feature you are attempting to use is not supported with the version of Storage Mirroring you are using.
-171	The command you are attempting to use is not supported by the Storage Mirroring version you are using.
-172	There is an error with the log file on the target.
-173	The target disk is full.
-174	The target disk has sufficient disk space.
-175	There was an error reading from or writing to the log file.
-176	The memory-based queue is currently being used to store data.
-177	The disk-based queue is currently being used to store data.
-178	A restore is required.
-179	The ID the service supplied by the driver is invalid.
-180	The path you are attempting to block is the child of a path that is currently blocked.
-181	The path you are attempting to block is the parent of a path that is currently blocked.
-182	Target path blocking is disabled.
-183	The connection ID specified for the block or unblock command is invalid.
-184	No command objects were found in the queue.
-185	The target is discarding operations from the target queue.
-186	The target is not discarding operations from the target queue.
-201	The monitor you are attempting to create already exists.
-202	The address specified is not configured to be monitored.
-203	The monitor you are attempting to configure already exists.
-204	The monitor you are attempting to disable does not exist.
-205	The monitor set you are attempting to enable is already in use.
-206	The monitor you are attempting to use with monitor set does not exist.
-209	The machine you are attempting to failback is not failed over.
-210	The machine you are attempting to failover is already failed over.
-211	The machine you are attempting to failover is not currently being monitored.
-213	IP address placeholders on the target are unavailable.

Error Code	Description
-214	The specified target NIC was not found. Failover monitoring is not started.
-215	The Storage Mirroring source module, which contains the failover module, is not loaded.
-216	The source was unable to set the state of the replica.
-217	The source is not responding to an ICMP ping. If the source is being monitored, it will failover.
-2201	There was an error communicating with the e-mail server.
-2202	There was an error connecting with the e-mail server.
-2203	E-mail notification is disabled.
-2204	E-mail notification is enabled.
-2205	E-mail notification will not work without Internet Explorer version 5.0 and WMI installed.
-2206	E-mail notification will not work without Internet Explorer version 5.0 installed.
-2207	There was an error sending an e-mail.
-2208	There was an error sending a test e-mail.
-2209	There was a WMI error connecting to the e-mail server
-2210	E-mail notification will not work without WMI installed.
-2211	The Event Viewer settings for e-mail notification do not match the Storage Mirroring recommended settings.
-2212	The e-mail notification setting specified is incorrect.
-2213	The e-mail address already exists in the notification list.
-2214	The e-mail ID specified is invalid.
-2215	The e-mail format specified is invalid.
-2216	The e-mail address specified is not in the notification list.
-2217	The e-mail address notification list is empty.
-2218	The e-mail warning setting is not set.
-2219	The e-mail test warning setting is not set.
2200	E-mail notification is functioning properly.

DTCL Scripting Examples

This section shows examples of different script files.

Using batch files

To use a Storage Mirroring batch file, you must first create the following one-line batch file which initiates the Storage Mirroring Command Line Client and specifies the name of the script to execute.

Storage Mirroring.BAT

```
cd C:\Program Files\OpenView\Storage Mirroring
cmd /c DTCL -f "C:\Program Files\OpenView\Storage Mirroring\scriptname.txt"
```

Storage Mirroring.BAT Command	Description
cd C:\Program Files\OpenView\Storage Mirroring	Change to the directory where the Storage Mirroring program files are located.
cmd /c DTCL -f "C:\Program Files\OpenView\Storage Mirroring\scriptname.txt"	Starts an instance of the command prompt, carries out the specified command, and then stops. In this case, the specified command is to run the Storage Mirroring Command Line client with the file name option. The location of the script is specified after the file name option.

NOTE: If you do not specify a path for the script file, Storage Mirroring will look in the directory where the `DTCL -f` command was executed.

Basic getting started script

The following script will start a Storage Mirroring connection by creating a repset called Exchange and connecting it to the machine Jersey.

```
source indy;
repset create Exchange;
repset rule add c:\exchsrvr include, recursive;
repset rule add c:\exchsrvr\bin exclude;
repset rule add d:\exchsrvr include, recursive;
repset save;
connect Exchange to jersey map exact;
```

Basic getting started script with failover

The following script will start a Storage Mirroring connection by creating a repset called Exchange and connecting it to the machine Jersey. This script will also set up failover.

```
source indy;
repset create Exchange;
repset rule add c:\exchsrvr include, recursive;
repset rule add c:\exchsrvr\bin exclude;
repset rule add d:\exchsrvr include, recursive;
repset save;
connect Exchange to jersey map exact;
target jersey;
monitor create indy;
monitor move 205.31.4.193 to nic 3 interval 5 timeout 25;
monitor start indy;
```

Basic restore script

The following script will restore the data in the Exchange replication set from the target machine to the source machine.

```
source indy;  
restore Exchange to jersey nooverwritenewer, usetargetdb;
```

Basic getting started script with variables

The following script uses variables to identify the source, target, and replication set. These variables are then used in the DTCL commands. The \$connectionID variable will contain the connection ID for the connection established. This variable can then be used to establish a verification schedule.

```
$TheTarget = "jersey";  
$TheSource = "indy";  
$TheRepset = "Exchange";  
source $TheSource;  
repset create $TheRepset;  
repset rule add c:\exchsrvr include, recursive;  
repset rule add c:\exchsrvr\bin exclude;  
repset rule add d:\exchsrvr include, recursive;  
repset save;  
$ConnectionID = connect $TheRepset to $TheTarget map exact;  
verify $ConnectionID sync, newer, checksum, every 12 hr;
```

Basic flow control script

The following script uses the IF conditional, FOR loop, and WHILE loop commands. In the following examples the FOR loop will start a mirror for all connections between 1 and 10. The IF conditional will return the error command "mirror failed to start" if a value of 0 is not obtained. The WHILE loop will wait for a mirror to end and disconnect as long as the returned value is less than 10. The IF conditional will return the error message "failed to disconnect" if a value of 0 is not obtained.

```
#Start mirror for all connections 1-10. #  
#Notify user of any errors. #  
SOURCE indy;  
FOR $x = 1 TO 10 DO  
    $ret = MIRROR START $x;  
    IF $ret != 0 THEN  
        WRITE "Mirror failed to start";  
        WRITE $x;  
        WRITE $ret;  
    END  
END  
  
#Wait for mirror to end and disconnect all connections 1-10. #  
#Notify user of any errors. #  
SOURCE indy;  
$x = 1;  
WHILE $x <= 10 DO  
    waitonmirror $x;  
    $ret = disconnect $x;  
    IF $ret != 0 THEN  
        WRITE "Failed to disconnect";  
        WRITE $x;  
        WRITE $ret;  
    END  
    $x = $x + 1;  
END
```

B ► Application Failover

Storage Mirroring simulates application failover by replicating and failing over key application functions like user data, shares, machine names, IP addresses, and so on. Functions like service availability can also be triggered during failover by using the Storage Mirroring flexible scripting features. These scripts allow you to create and execute scripts on the target machine before and after failover and failback. You can also create and execute scripts on the source machine after failback. These scripts may contain any valid Windows command, executable, or batch file

Examples of functions specified in pre- and post-failover and failback scripts include:

- ◆ Stopping services on the target before failover because they may not be necessary while the target is standing in for the source
- ◆ Stopping services on the target that need to be restarted with the source's machine name and IP address
- ◆ Starting services or loading or applications that are in an idle, standby mode waiting for failover to occur
- ◆ Notifying the administrator before and after failover or failback occurs
- ◆ Stopping services on the target after failback because they are no longer needed
- ◆ Stopping services on the target that need to be restarted with the target machine's original name and IP address

With these flexible scripting features, application failover using Storage Mirroring can be seamless to the end user. HP tests many of the popular applications on the market today including Exchange, SQL, Oracle®, and Lotus Domino®, just to name a few. The results of these testing procedures are written up into formal Application Notes that describe how Storage Mirroring should be configured to work correctly with certain applications.

Script Sample

The following batch file might be used to stop and start Microsoft Exchange services on a Windows machine. This file is just a sample to demonstrate the kinds of commands and functions you can perform using scripts. You will need to modify this script to work in your environment and to achieve the desired outcome.

SAMPLE_NT_POSTOVER.BAT

```
rem The following command sends a message to the network administrator
rem You need to insert the network administrator's computer name for
rem <computername>
net send <computername> Source has failed. Failover is starting.

rem These commands start the initial services to run the patch
Net Start "Microsoft Exchange System Attendant"
Net Start "Microsoft Exchange Directory"

rem These commands patch the database
cd \
cd exchsrvr
cd bin
isinteg -patch

rem This command stops the services
Net Stop "Microsoft Exchange System Attendant" /y

rem These commands start the services
Net Start "Microsoft Exchange System Attendant"
Net Start "Microsoft Exchange Directory"
Net Start "Microsoft Exchange Message Transfer Agent"
Net Start "Microsoft Exchange Information Store"
Net Start "Microsoft Exchange Event Service"

rem The next command is for Exchange Version 5.0 only
rem Remove the rem characters on the last line if you want this line to
rem execute
rem Net Start "Microsoft Exchange Directory Synchronization"
```


C ► Disaster Recovery for MSCS

A Microsoft cluster is a group of independent servers, called nodes, working together as a unit. Clients interact with a cluster as though it is a single server. When a node in the cluster fails, Microsoft Cluster Server (MSCS) responds by moving the work from the failed node to the other node in the cluster. Microsoft documents and recommends a standard cluster configuration containing two machines sharing a SCSI disk. In this configuration, a single copy of the data resides on the SCSI disk and is shared between the cluster nodes.

Storage Mirroring works in conjunction with MSCS to provide disaster recovery for an MSCS cluster. Storage Mirroring continues to protect data when a source or target drive is moved from one node of a cluster to another either due to administrator intervention or a node failure.

NOTE:	Each cluster's private network should be a unique subnet so that the Management Console will not attempt to use an unreachable private network..
--------------	--

Cluster Configurations

Storage Mirroring flexibility allows multiple configurations to be used with a cluster. Review the brief clustering overview on the following page and then use the appropriate instructions based on your specific configuration.


- ◆ [Cluster to Cluster Configuration](#) on page C-5
- ◆ [Cluster to Standalone Configuration](#) on page C-13
- ◆ [Standalone to Cluster Configuration](#) on page C-21

NOTE: The first two configurations where the source is a cluster require the use of the Double-Take Source Connection resource on the cluster. Without it, there is no mechanism to bring the Storage Mirroring connection online when there is a change in node ownership. When there is a failure or the administrator manually changes node ownership, Storage Mirroring will disconnect the connection on the failed node before cluster administrator moves control to the new node. When the resources are brought online on the new node, the Double-Take Source Connection resource establishes a new Storage Mirroring connection on the now active node.

MSCS and Storage Mirroring

The Microsoft Cluster Server configuration typically contains two machines, called nodes, sharing at least one SCSI disk. Microsoft recommends at least two network interfaces where one or more is dedicated to node to node traffic. The cluster contains resources such as services and devices that are managed by the cluster and allowed to move from one node to the other. The shared SCSI disk is considered a resource as well as drive shares, services, applications, IP addresses, computer names, and so on. Resources can be owned by only one node of the cluster at a time. If the owning cluster node fails, the second node takes ownership and control of the resources.


Resource groups provide an administrative method for collecting multiple resources together for simple management. When a resource group moves from one node to the other, all of the resources within that group move as well.



By adding the Double-Take Source Connection resource to this configuration, you can provide disaster recovery for your cluster. The Source Connection resource allows you to specify basic Storage Mirroring connection parameters. By defining a replication set and a Source Connection resource per cluster group, you can replicate your entire cluster or specific cluster groups to a standalone machine or to the node of another cluster.

Providing Disaster Recovery Using Storage Mirroring

According to Microsoft specifications, each application or network service running on a cluster should be in a separate group. A group should only contain multiple applications if the applications are dependent on each other for proper execution. Each group should also contain a unique network name, IP address, and its own shared disk, in order to present a virtual server to the network clients. To provide disaster recovery, you need to protect the data that is contained in each group. Storage Mirroring can provide that protection.



Cluster to Cluster Configuration

The following instructions are for a cluster to cluster configuration.

Configuring the first cluster

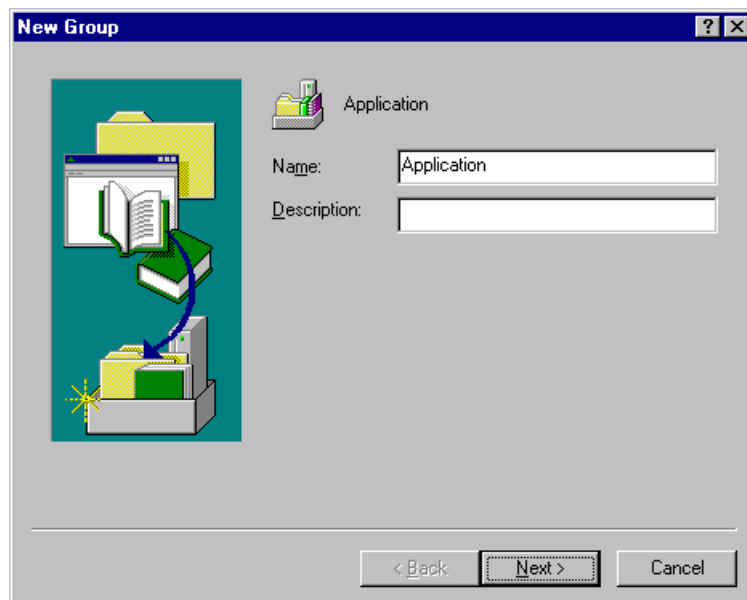
For the first cluster you will need to configure the virtual server, install and configure Storage Mirroring and configure the Storage Mirroring replication set.

Configure the virtual server on the first cluster

Client machines need to access a server when looking for an application, printer, file share, or other resource. Network clients will not be aware which node owns the resource. Therefore, a virtual server needs to be created. This virtual server must have an IP address, a network name, and a disk resource. The virtual server will move between nodes as necessary, transparent to the clients. The clients will always be able to see and access the resources hosted on the virtual server, no matter which node owns it.

To create a virtual server you must create a resource group and add to it an IP address, network name, and a physical disk.

- I. Create a resource group for the virtual server.
 - a. Select **Start, Programs, Administrative Tools, Cluster Administrator** and connect to the first cluster.
 - b. Right-click the **Groups** folder on the left pane of the Cluster Administrator and select **New, Group**.
 - c. Specify the **Name** and **Description** and click **Next** to continue.



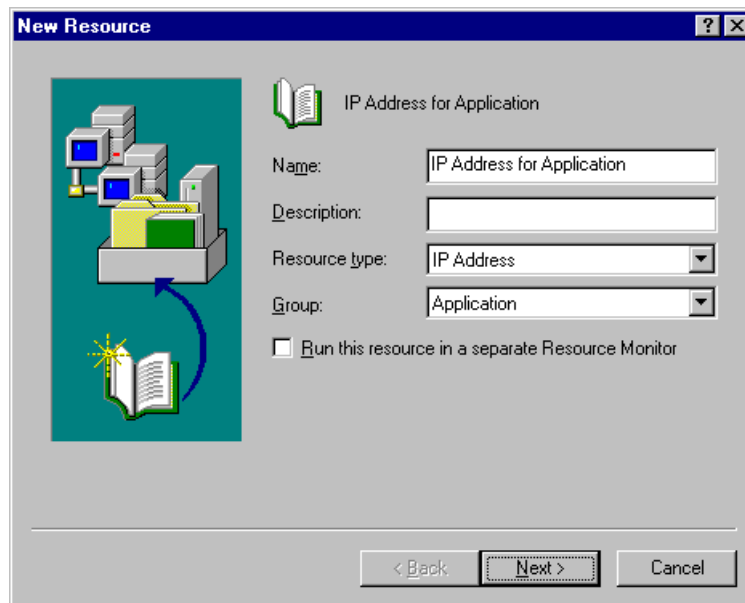
- d. Specify zero or one cluster node as the **Preferred Owner**. Do not specify both nodes. If a preferred node is selected for a resource group, then the group will automatically move to its preferred node when that node comes online. This does not affect manual moves of the group when both nodes are online.

NOTE: For the Preferred Node functionality to work for a resource group, Failback must be also be enabled on the Properties Failback tab of that resource group.

- e. Click **Finish** to complete the creation of the new group.

NOTE: You will be notified that the group was created successfully. Click **OK** to acknowledge the message and return to the Cluster Administrator main screen.

2. Create an IP Address resource for the new group.
 - a. Right-click the group that you just created and select **New, Resource**.
 - b. Specify the following fields on the New Resource dialog box:

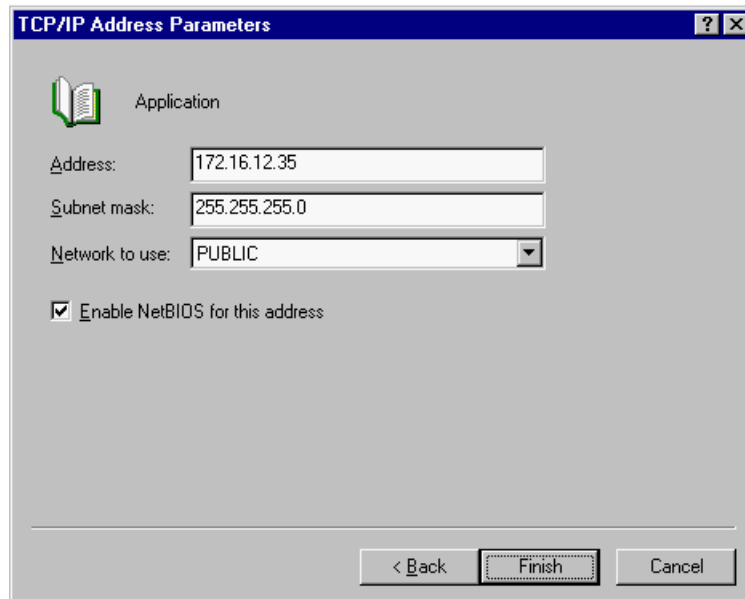


The 'New Resource' dialog box is shown. It has a title bar with a question mark and a close button. On the left is an icon of a server rack with a blue arrow pointing to a book icon. The main area is titled 'IP Address for Application'. It contains the following fields:

- Name:** A text box containing 'IP Address for Application'.
- Description:** An empty text box.
- Resource type:** A dropdown menu with 'IP Address' selected.
- Group:** A dropdown menu with 'Application' selected.
- ☐ **Run this resource in a separate Resource Monitor**

At the bottom are three buttons: '< Back', 'Next >', and 'Cancel'.

- ◆ **Name**—Specify a name that identifies this resource as the IP address for the group. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **IP Address**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
 - d. Verify that both nodes appear as the **Possible Owners** and click **Next** to continue.
 - e. There are no **Dependencies** required for an IP address resource. Click **Next** to continue.
 - f. Specify the following fields on the TCP/IP Address Parameters dialog box:



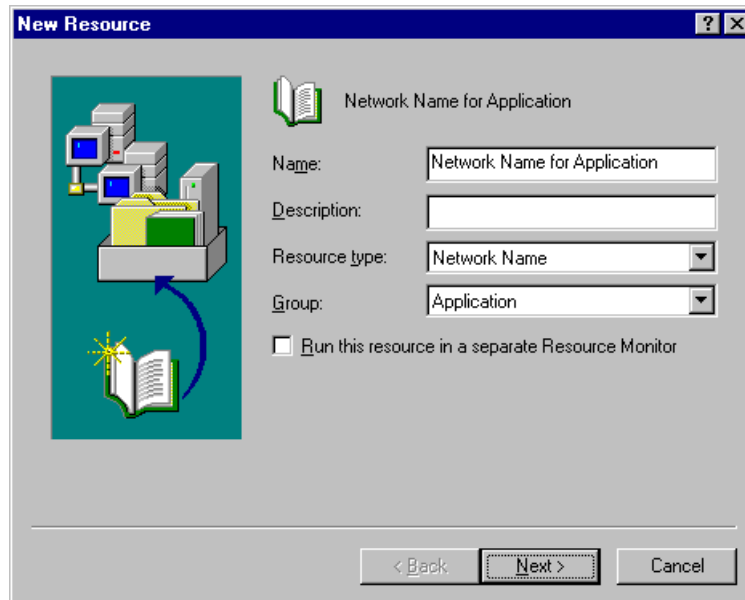
The 'TCP/IP Address Parameters' dialog box is shown. It has a title bar with a question mark and a close button. On the left is an icon of a book. The main area is titled 'Application'. It contains the following fields:

- Address:** A text box containing '172.16.12.35'.
- Subnet mask:** A text box containing '255.255.255.0'.
- Network to use:** A dropdown menu with 'PUBLIC' selected.
- ☒ **Enable NetBIOS for this address**

At the bottom are three buttons: '< Back', 'Finish', and 'Cancel'.

- ◆ **Address**—Enter the static IP address that will be assigned to the virtual server.
 - ◆ **Subnet mask**—Enter the subnet mask associated with the IP address you just entered.
 - ◆ **Network to use**—If you have more than one route for network traffic defined, specify the network that this IP address will use. This should be the same network that the clients are using.
- g. Click **Finish** to complete the creation of the IP Address resource.

3. Create a Network Name resource for the new group.
 - a. Right-click the group and select **New, Resource**.
 - b. Specify the following fields on the New Resource dialog box:

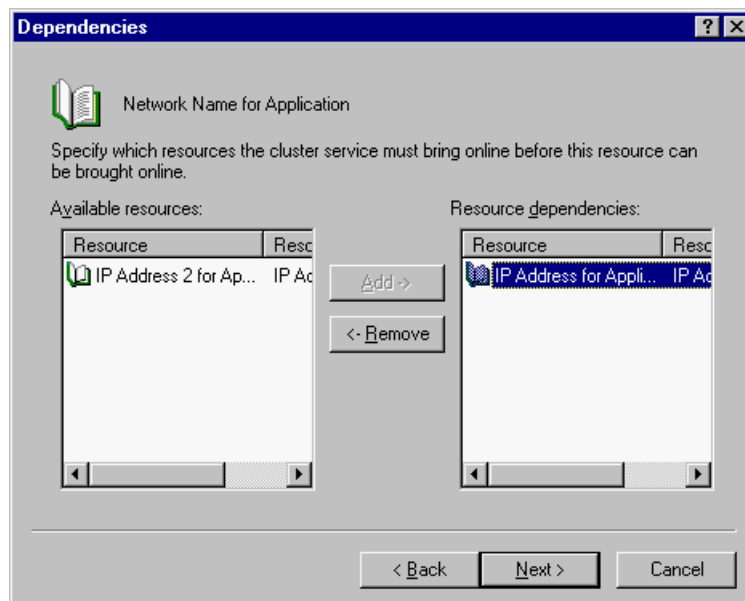


The "New Resource" dialog box is shown. It has a title bar with a question mark and a close button. On the left is an icon of a server rack with a folder and a book. The main area is titled "Network Name for Application". It contains the following fields:

- Name:** A text box containing "Network Name for Application".
- Description:** An empty text box.
- Resource type:** A dropdown menu with "Network Name" selected.
- Group:** A dropdown menu with "Application" selected.
- ☐ **Run this resource in a separate Resource Monitor**

At the bottom are three buttons: "< Back", "Next >", and "Cancel".

- ◆ **Name**—Specify a name that identifies this resource as the virtual server name. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Network Name**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
 - d. Verify that both nodes appear as **Possible Owners** and click **Next** to continue.
 - e. An IP address must be present in order for a network name to be assigned. Therefore, move the IP address resource associated with the virtual server to the **Resource dependencies** list and click **Next** to continue.



The "Dependencies" dialog box is shown. It has a title bar with a question mark and a close button. On the left is an icon of a book. The main area is titled "Network Name for Application". Below the title is the text: "Specify which resources the cluster service must bring online before this resource can be brought online."

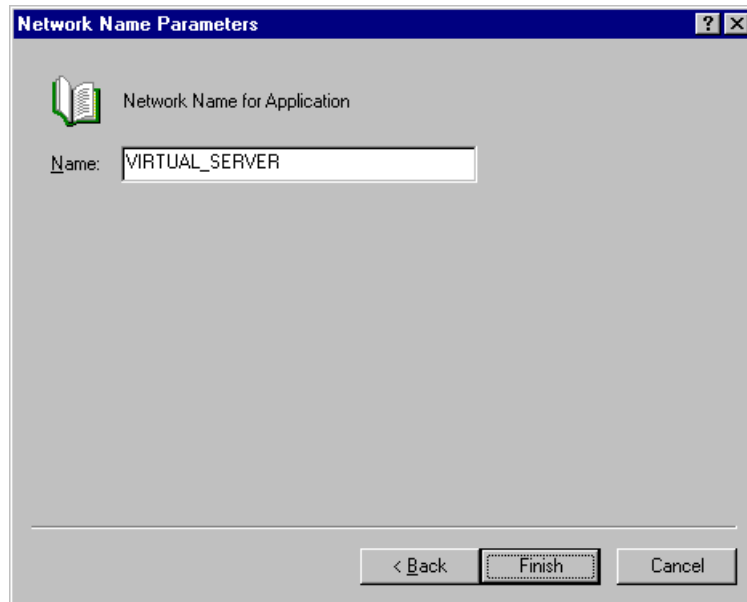
There are two lists:

- Available resources:** A list box containing one item: "IP Address 2 for Ap... IP Ac".
- Resource dependencies:** A list box containing one item: "IP Address for Appli... IP Ac".

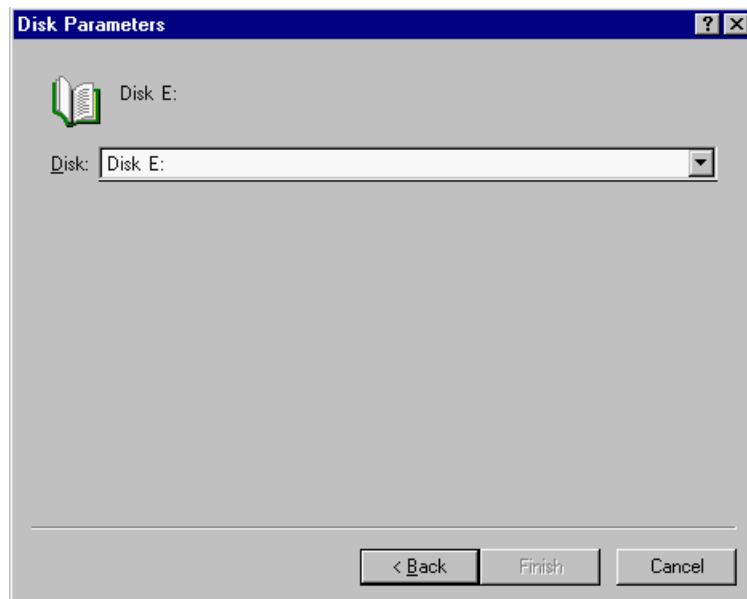
Between the two lists are two buttons: "Add >" and "< Remove".

At the bottom are three buttons: "< Back", "Next >", and "Cancel".

- f. Specify the Network Name Parameters by entering the virtual name of the server. This is the name that clients will look for on the network and the name that will register with DNS. It must be unique both on the cluster and on the network.



- g. Click **Finish** to complete the creation of the Network Name resource.
4. Create a Physical Disk resource for the new group.
- a. Right-click the group and select **New, Resource**.
- b. Specify the following fields on the New Resource dialog box:
- ◆ **Name**—Specify a name that identifies the disk drive associated with the virtual server. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Physical Disk**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
- d. Verify that both nodes appear as **Possible Owners** and click **Next** to continue.
- e. No other resources are required as dependencies for disk resources. Click **Next** to continue.
- f. Specify the disk drive associated with the physical disk resource.



- g. Click **Finish** to complete the creation of the Physical Disk resource.

The application group now contains all of the resources necessary to identify a virtual server. All resources should be offline at this time. If any are online, they need to be taken offline by right-clicking the resource and selecting **Take Offline**.

Install and Configure Storage Mirroring on the first cluster

1. If you have not already done so, install Storage Mirroring on the first node of the first cluster using the installation defaults.

NOTE: The Storage Mirroring installation will prompt you to reboot the machine. You can reboot at a later time, but in order for Storage Mirroring to function properly and for you to finish these instructions, you must reboot.

For complete installation details, see the *Getting Started* guide.

2. Disable Storage Mirroring's standard connection controls so that the Double-Take Source Connection resource, that you will be configuring later, can control the Storage Mirroring connections.
 - a. Start the Storage Mirroring Management Console by selecting **Start, Programs, Storage Mirroring, Management Console**.
 - b. Right-click the first node of the source cluster and select **Properties**.
 - c. Select the **Setup** tab.
 - d. By default, the **Automatically Reconnect During Source Initialization** check box will be selected. Disable this option by clearing the check box.
 - e. Click **OK** to save the changes.
3. Repeat steps 1 and 2 on the second node of the first cluster.

Configure the Storage Mirroring replication set

In order for the clusters to be synchronized, the data that is changed on the source cluster must be replicated to the target. Storage Mirroring handles this task by establishing a replication set which identifies the data that is changing.

1. On the node owning the group you wish to protect, start the Storage Mirroring Management Console by selecting **Start, Programs, Storage Mirroring, Management Console**.
2. Double-click on the node on the left pane of the Management Console to login.
3. Right-click the node and select **New, Replication Set**.
4. Enter a name for the replication set and press **Enter**.
5. On the right pane of the Management Console, select the directory and files in the tree that are associated with the group and virtual server you are protecting.
6. Right-click the replication set name and select **Save**.
7. Right-click the replication set that you just created and select **Properties**.
8. Record the exact drive and directories of each path displayed in the Replication Set Properties table provided. Place a check mark or X in the Include, Exclude, and Recurse Sub-directories columns to identify which parameters apply to the specified path.

Replication Set Properties

Drive and Directories	Include	Exclude	Recurse Sub-directories

9. Double-click on the second node on the left pane of the Management Console to login.
10. Right-click the node and select **New, Replication Set**.
11. Enter the exact, case-sensitive name for the replication set as specified on the first node and press **Enter**.

12. Because the second node does not currently own the files, you will not be able to browse to select the data like you did on the first node. Therefore, you will need to manually enter the replication set data that you recorded in the **Replication Set Properties** table.

- a. Right-click the replication set that you just created and select **Properties**.
- b. Click **Add**.
- c. Specify one of the drive and directory paths that you recorded in the table **Replication Set Properties** on page C-9. Be sure and mark the correct **Include**, **Exclude** and **Recurse sub-directories** options that need to be applied.

NOTE: Each replication set rule must be identical to the replication set rule on the first node in order for the disaster recovery process to work correctly.

- d. Click **OK** to save the replication set rule.
- e. Repeat steps b-d for each path and directory on the first node.

NOTE: Each drive and directory will appear in the Replication Set Properties even though the second node may not have access to these locations right now. That is not a problem.

13. Right-click the replication set name and select **Save**.

At this point, you are not finished with the configuration of the first cluster. The remaining steps for the first cluster are dependent on components of the second cluster. Therefore, you need to configure the second cluster and then you will come back and complete the remaining configuration steps on the first cluster.

Configuring the second cluster

For the second cluster you will need to configure the virtual server and bring it online.

Configure the virtual server on the second cluster

1. Repeat the steps in **Configure the virtual server on the first cluster** on page C-5 to create another, unique virtual server but this time on the second cluster.
2. Bring the virtual server and its resources online.

Install Storage Mirroring on the second cluster

1. If you have not already done so, install Storage Mirroring on the first node of the second cluster using the installation defaults.

NOTE: The Storage Mirroring installation will prompt you to reboot the machine. You can reboot at a later time, but in order for Storage Mirroring to function properly and for you to finish these instructions, you must reboot.

For complete installation details, see the *Getting Started* guide.

2. Disable Storage Mirroring's standard connection controls so that the Double-Take Source Connection resource can control the Storage Mirroring connections.
 - a. Start the Storage Mirroring Management Console by selecting **Start, Programs, Storage Mirroring, Management Console**.
 - b. Right-click the target server and select **Properties**.
 - c. Select the **Setup** tab.
 - d. By default, the **Automatically Reconnect During Source Initialization** check box will be selected. Disable this option by clearing the check box.
 - e. Click **OK** to save the changes.
3. Repeat steps 1 and 2 on the second node of the second cluster.

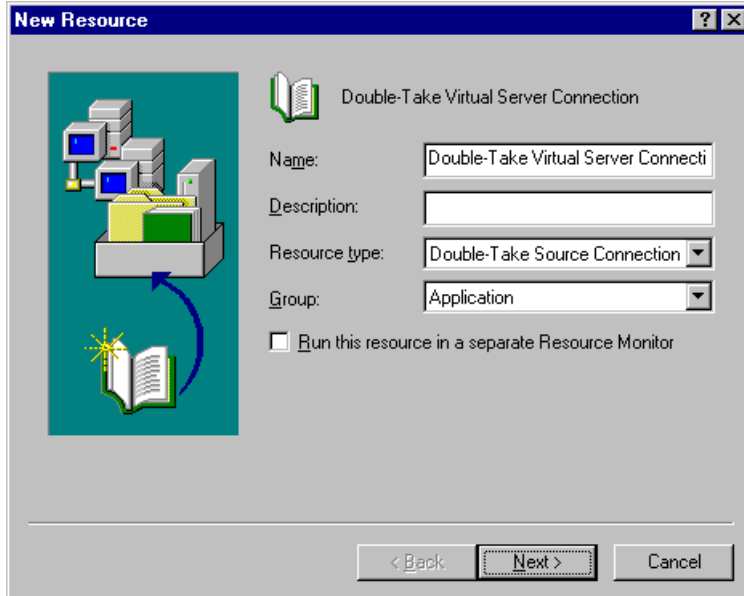
Completing the first cluster configuration

To finalize the configuration, you need to create a Double-Take Source Connection resource and bring both the virtual server and Source Connection components online. Before beginning, make sure that the resource group and all of its resources are offline on the second cluster.

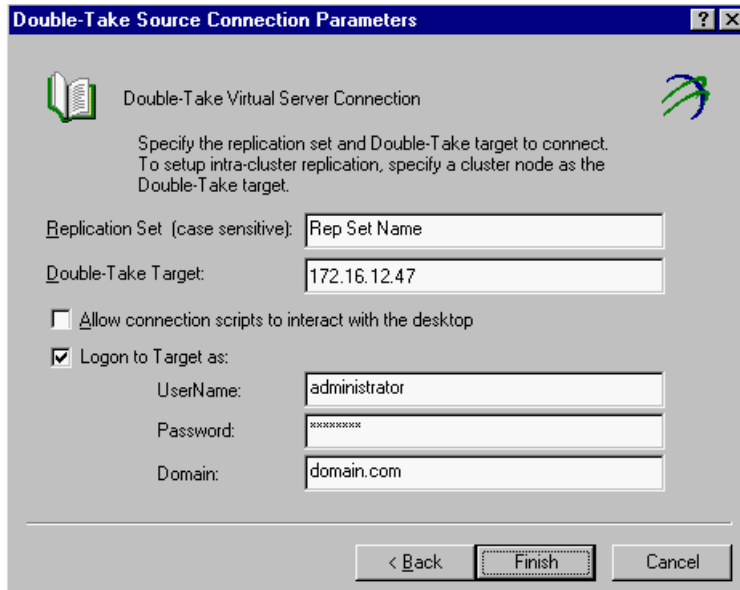
Configure the Double-Take Source Connection resource

The Double-Take Source Connection resource controls the Storage Mirroring connections. You need to configure this resource through the Cluster Administrator.

1. Select **Start, Programs, Administrative Tools, Cluster Administrator**.
2. Right-click the application group you were working with earlier and select **New, Resource**.
3. Specify the following fields on the New Resource dialog box:



- ◆ **Name**—Specify a name that indicates this is the Storage Mirroring virtual server connection.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Double-Take Source Connection**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
4. Click **Next** to continue.
 5. Verify that both nodes appear as **Possible Owners** and click **Next** to continue.
 6. To keep the Double-Take Source Connection resource from coming online before the physical disk, make this resource dependent on the physical disk resource created earlier in [Configure the virtual server on the first cluster](#). Click **Next** to continue.
 7. Specify the following on the Double-Take Source Connection Parameters dialog box:



-
- ◆ **Replication Set**—Specify the name of the Storage Mirroring replication set. This name is case-sensitive and should be the same name as specified in [Configure the Storage Mirroring replication set](#) on page C-9.
 - ◆ **Double-Take Target**—Specify the IP address of the virtual server of the second cluster. This is the IP address resource created in [Configure the virtual server on the second cluster](#) on page C-10.
 - ◆ **Allow connection scripts to interact with the desktop**—Mark this check box if you want to display the connection information in a command prompt dialog box.

NOTE: If the script is set to interact with the desktop, the results will be displayed on the owning node. The resource will remain at the **Online pending** status until the script is manually acknowledged by a user. Additionally, if you administer the cluster through Terminal Services, the command prompt windows will not appear on your Terminal Services client, only on the nodes themselves. Therefore, this option should not be a permanent setting, but should be used for troubleshooting if you are having problems establishing a connection.

For additional information on the scripts used, see [Understanding the Double-Take Source Connection resource](#) on page C-27.

-
- ◆ **Logon to Target as**—Allows you to specify an ID, password and domain to be used when logging on to the target.

NOTE: The user must be a member of the **Double-Take Admin** security group on all nodes of the cluster running a Double-Take Source Connection resource.

If you do not specify a logon, the ID used to start the cluster service will be used. Verify that the ID has **Double-Take Admin** access.

-
8. Click **Finish** to complete the creation of the Double-Take Source Connection resource.

Bring the Components Online

From the Cluster Administrator (**Start, Programs, Administrative Tools, Cluster Administrator**) bring the Double-Take Source Connection resource and the virtual server resources online.

Notice the warning signs are removed from the resource and the group. The **State** column indicates the resource is online and functioning properly. The data is now being replicated from the source cluster to the target cluster.

NOTE: If the State column indicates failed, as it did prior to bringing the resource online, the connection has not executed properly. When the resource has failed, the Storage Mirroring connection is not active and your data is not being protected in the event of a disaster.

To troubleshoot a failed Source Connection resource, right-click the resource and select **Properties**. Select the **Parameters** tab. Mark the **Allow connection script to interact with desktop** checkbox and click **OK**. When you attempt to bring the resource online, the results from the connection script will be displayed on screen on the node running the script allowing you to see what is causing the failure.

When using the **Allow connection script to interact with desktop** option with the Double-Take Source Connection resource, the resource will report a pending state until the command prompt is closed. The actual status of the resource can be obtained by viewing the script output in the command prompt. Additionally, closing the command prompt by clicking on the close button in the upper right hand corner of the prompt will cause the resource to report a failure. Closing the command prompt by pressing a key on the keyboard ensures that the resource will report its correct state. Additionally, if you administer the cluster through Terminal Services, the command prompt windows will not appear on your Terminal Services client, only on the nodes themselves.

For additional information on the scripts used, see [Understanding the Double-Take Source Connection resource](#) on page C-27.

If you need to terminate the Storage Mirroring connection, right-click the resource and select **Take Offline**. The resource will appear, offline, in the Cluster Administrator. If the Double-Take Source Connection resource is offline, data is not being replicated and is not protected in the event of a disaster.

Cluster to Standalone Configuration

The following instructions are for a cluster to standalone server configuration.

Configuring the cluster

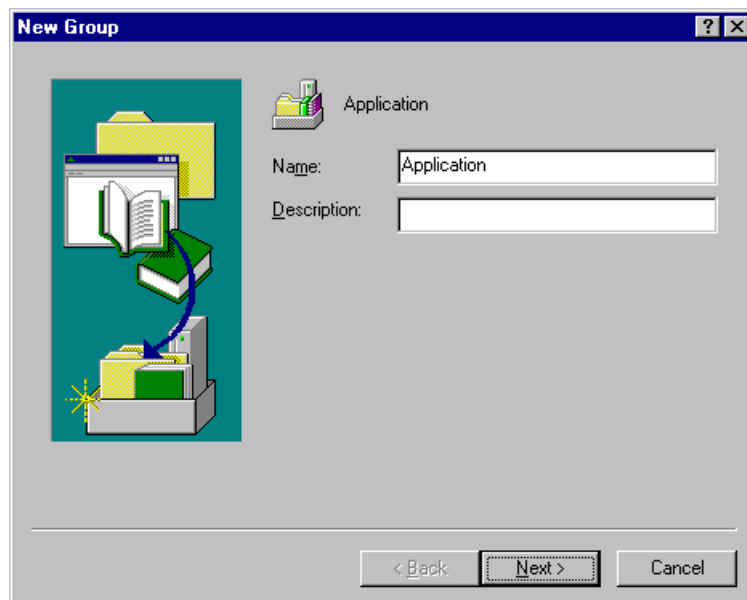
For the cluster you will need to configure the virtual server, install and configure Storage Mirroring, configure the Storage Mirroring replication set, and create the Double-Take Source Connection resource.

Configure the virtual server on the cluster

Client machines need to access a server when looking for an application, printer, file share, or other resource. Network clients will not be aware which nodes owns the resource. Therefore, a virtual server needs to be created. This virtual server must have an IP address, a network name, and a disk resource. The virtual server will move between nodes as necessary, transparent to the clients. The clients will always be able to see and access the resources hosted on the virtual server, no matter which node owns it.

To create a virtual server you must create a resource group and add to it an IP address, network name, and a physical disk.

- I. Create a resource group for the virtual server.
 - a. Select **Start, Programs, Administrative Tools, Cluster Administrator** and connect to the first cluster.
 - b. Right-click the **Groups** folder on the left pane of the Cluster Administrator and select **New, Group**.
 - c. Specify the **Name** and **Description** and click **Next** to continue.



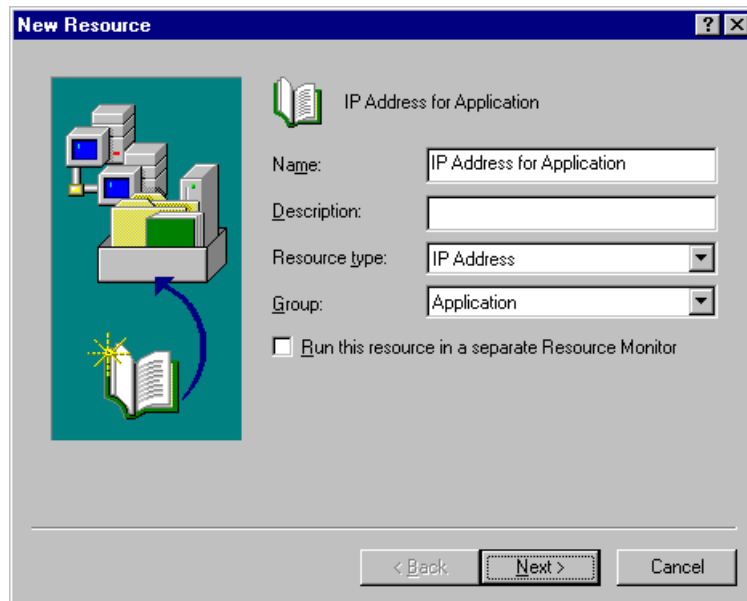
- d. Specify zero or one cluster node as the **Preferred Owner**. Do not specify both nodes. If a preferred node is selected for a resource group, then the group will automatically move to its preferred node when that node comes online. This does not affect manual moves of the group when both nodes are online.

NOTE: For the Preferred Node functionality to work for a resource group, Failback must be also be enabled on the Properties Failback tab of that resource group.

- e. Click **Finish** to complete the creation of the new group.

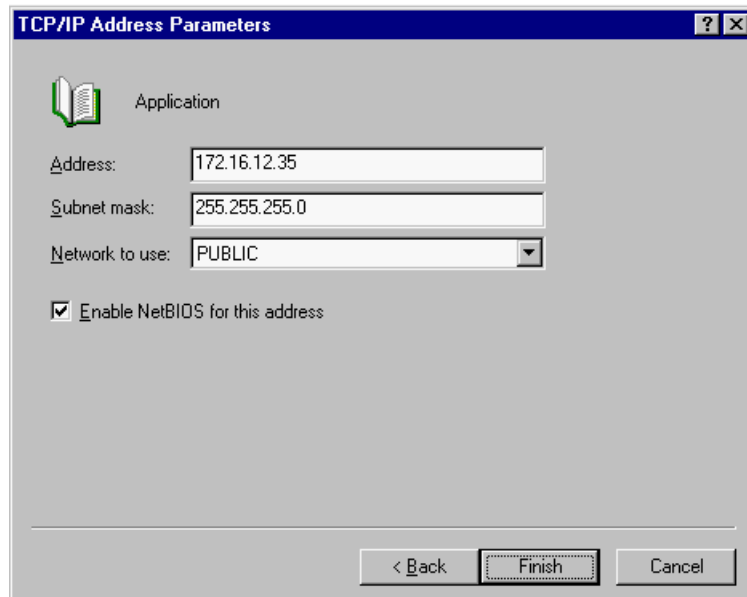
NOTE: You will be notified that the group was created successfully. Click **OK** to acknowledge the message and return to the Cluster Administrator main screen.

2. Create an IP Address resource for the new group.
 - a. Right-click the group that you just created and select **New, Resource**.
 - b. Specify the following fields on the New Resource dialog box:



The 'New Resource' dialog box is shown. It has a title bar with a question mark and a close button. On the left is an icon of a server rack with a blue arrow pointing to a book icon. On the right, the text 'IP Address for Application' is displayed. Below this, there are four fields: 'Name' with the value 'IP Address for Application', 'Description' (empty), 'Resource type' with a dropdown menu showing 'IP Address', and 'Group' with a dropdown menu showing 'Application'. At the bottom left is a checkbox labeled 'Run this resource in a separate Resource Monitor' which is unchecked. At the bottom right are three buttons: '< Back', 'Next >', and 'Cancel'.

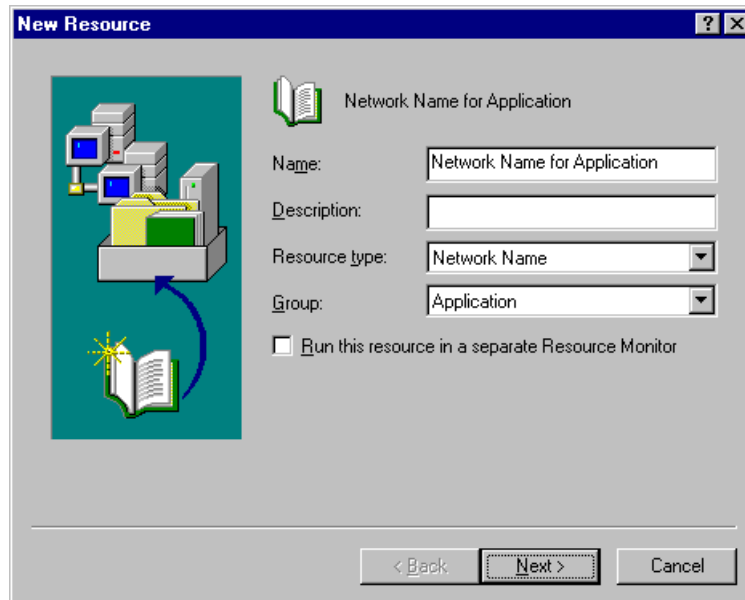
- ◆ **Name**—Specify a name that identifies this resource as the IP address for the group. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **IP Address**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
 - d. Verify that both nodes appear as the **Possible Owners** and click **Next** to continue.
 - e. There are no **Dependencies** required for an IP address resource. Click **Next** to continue.
 - f. Specify the following fields on the TCP/IP Address Parameters dialog box:



The 'TCP/IP Address Parameters' dialog box is shown. It has a title bar with a question mark and a close button. On the left is an icon of a book. On the right, the text 'Application' is displayed. Below this, there are three fields: 'Address' with the value '172.16.12.35', 'Subnet mask' with the value '255.255.255.0', and 'Network to use' with a dropdown menu showing 'PUBLIC'. At the bottom left is a checkbox labeled 'Enable NetBIOS for this address' which is checked. At the bottom right are three buttons: '< Back', 'Finish', and 'Cancel'.

- ◆ **Address**—Enter the static IP address that will be assigned to the virtual server.
 - ◆ **Subnet mask**—Enter the subnet mask associated with the IP address you just entered.
 - ◆ **Network to use**—If you have more than one route for network traffic defined, specify the network that this IP address will use. This should be the same network that the clients are using.
- g. Click **Finish** to complete the creation of the IP Address resource.

3. Create a Network Name resource for the new group.
 - a. Right-click the group and select **New, Resource**.
 - b. Specify the following fields on the New Resource dialog box:

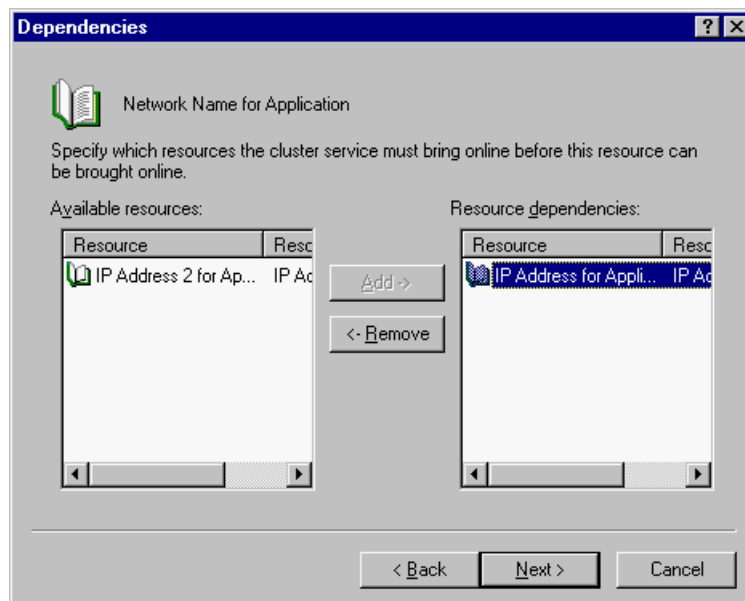


The 'New Resource' dialog box is shown. It has a title bar with a question mark and a close button. On the left is an icon of a server rack with a folder and a book. The main area is titled 'Network Name for Application'. It contains the following fields:

- Name: Network Name for Application
- Description: (empty)
- Resource type: Network Name (dropdown)
- Group: Application (dropdown)
- ☐ Run this resource in a separate Resource Monitor

At the bottom are three buttons: '< Back', 'Next >', and 'Cancel'.

- ◆ **Name**—Specify a name that identifies this resource as the virtual server name. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Network Name**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
 - d. Verify that both nodes appear as **Possible Owners** and click **Next** to continue.
 - e. An IP address must be present in order for a network name to be assigned. Therefore, move the IP address resource associated with the virtual server to the **Resource dependencies** list and click **Next** to continue.

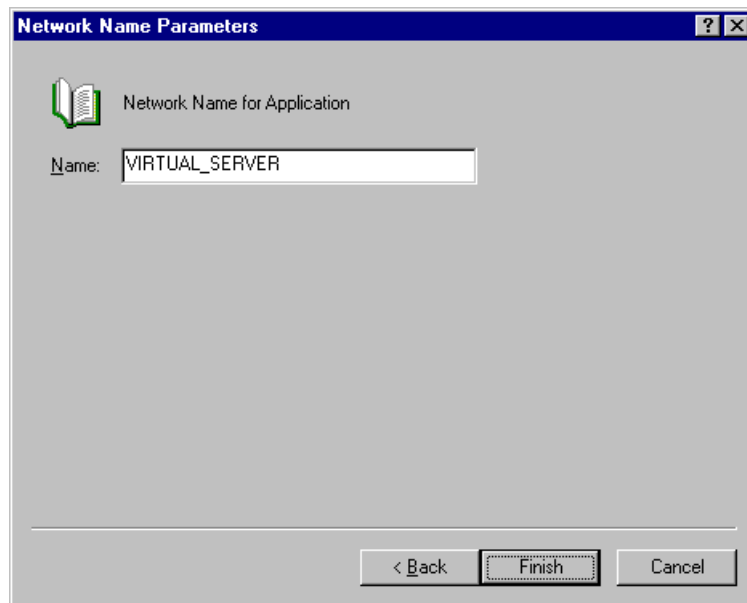


The 'Dependencies' dialog box is shown. It has a title bar with a question mark and a close button. On the left is an icon of a book. The main area is titled 'Network Name for Application'. Below the title is the text: 'Specify which resources the cluster service must bring online before this resource can be brought online.' There are two lists:

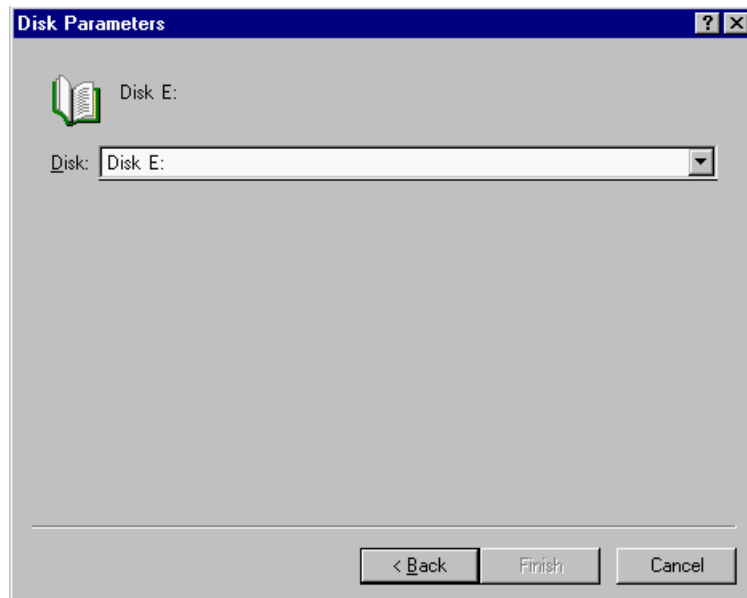
- Available resources:** A table with two columns: 'Resource' and 'Resc'. It contains one entry: 'IP Address 2 for Ap...' with 'IP Ac' in the 'Resc' column.
- Resource dependencies:** A table with two columns: 'Resource' and 'Resc'. It contains one entry: 'IP Address for Appli...' with 'IP Ac' in the 'Resc' column.

Between the two lists are two buttons: 'Add >' and '< Remove'. At the bottom are three buttons: '< Back', 'Next >', and 'Cancel'.

- f. Specify the Network Name Parameters by entering the virtual name of the server. This is the name that clients will look for on the network and the name that will register with DNS. It must be unique both on the cluster and on the network.



- g. Click **Finish** to complete the creation of the Network Name resource.
4. Create a Physical Disk resource for the new group.
- a. Right-click the group and select **New, Resource**.
- b. Specify the following fields on the New Resource dialog box:
- ◆ **Name**—Specify a name that identifies the disk drive associated with the virtual server. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Physical Disk**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
- d. Verify that both nodes appear as **Possible Owners** and click **Next** to continue.
- e. No other resources are required as dependencies for disk resources. Click **Next** to continue.
- f. Specify the disk drive associated with the physical disk resource.



- g. Click **Finish** to complete the creation of the Physical Disk resource.

The application group now contains all of the resources necessary to identify a virtual server. All resources should be offline at this time. If any are online, they need to be taken offline by right-clicking the resource and selecting **Take Offline**.

Install and Configure Storage Mirroring on the cluster

1. If you have not already done so, install Storage Mirroring on the first node of the cluster using the installation defaults.

NOTE: The Storage Mirroring installation will prompt you to reboot the machine. You can reboot at a later time, but in order for Storage Mirroring to function properly and for you to finish these instructions, you must reboot.

For complete installation details, see the *Getting Started* guide.

2. Disable Storage Mirroring's standard connection controls so that the Double-Take Source Connection resource, that you will be configuring later, can control the Storage Mirroring connections.
 - a. Start the Storage Mirroring Management Console by selecting **Start, Programs, Storage Mirroring, Management Console**.
 - b. Right-click the first node of the cluster and select **Properties**.
 - c. Select the **Setup** tab.
 - d. By default, the **Automatically Reconnect During Source Initialization** check box will be selected. Disable this option by clearing the check box.
 - e. Click **OK** to save the changes.
3. Repeat steps 1 and 2 on the second node of the cluster.

Configure the Storage Mirroring replication set

In order for the clusters to be synchronized, the data that is changed on the cluster must be replicated to the target. Storage Mirroring handles this task by establishing a replication set which identifies the data that is changing.

1. On the node owning the group you wish to protect, start the Storage Mirroring Management Console by selecting **Start, Programs, Storage Mirroring, Management Console**.
2. Double-click on the node on the left pane of the Management Console to login.
3. Right-click the node and select **New, Replication Set**.
4. Enter a name for the replication set and press **Enter**.
5. On the right pane of the Management Console, select the directory and files in the tree that are associated with the group and virtual server you are protecting.
6. Right-click the replication set name and select **Save**.
7. Right-click the replication set that you just created and select **Properties**.
8. Record the exact drive and directories of each path displayed in the Replication Set Properties table provided. Place a check mark or X in the Include, Exclude, and Recurse Sub-directories columns to identify which parameters apply to the specified path.

Replication Set Properties

Drive and Directories	Include	Exclude	Recurse Sub-directories

9. Double-click on the second node on the left pane of the Management Console to login.
10. Right-click the node and select **New, Replication Set**.
11. Enter the exact, case-sensitive name for the replication set as specified on the first node and press **Enter**.

12. Because the second node does not currently own the files, you will not be able to browse to select the data like you did on the first node. Therefore, you will need to manually enter the replication set data that you recorded in the **Replication Set Properties** table.

- Right-click the replication set that you just created and select **Properties**.
- Click **Add**.
- Specify one of the drive and directory paths that you recorded in the table **Replication Set Properties** on page C-17. Be sure and mark the correct **Include**, **Exclude** and **Recurse sub-directories** options that need to be applied.

NOTE: Each replication set rule must be identical to the replication set rule on the first node in order for the disaster recovery process to work correctly.

- Click **OK** to save the replication set rule.
- Repeat steps b-d for each path and directory on the first node.

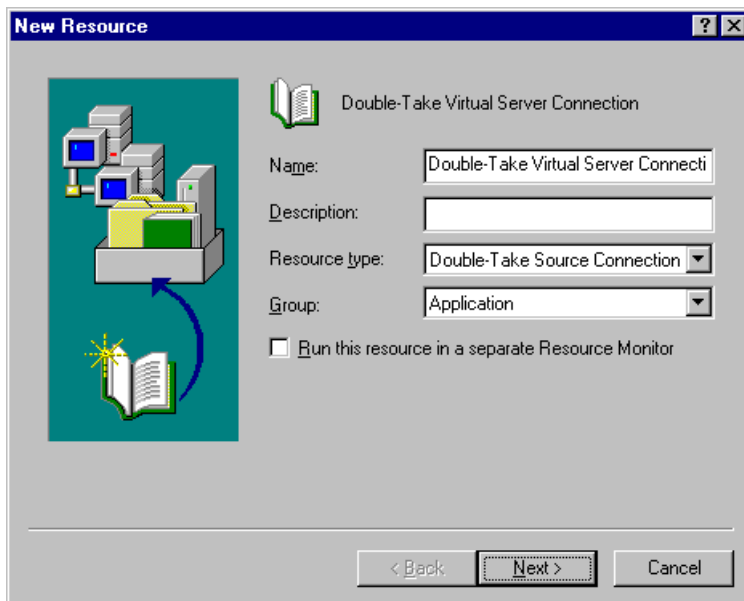
NOTE: Each drive and directory will appear in the Replication Set Properties even though the second node may not have access to these locations right now. That is not a problem.

13. Right-click the replication set name and select **Save**.

Configure the Double-Take Source Connection resource

The Double-Take Source Connection resource controls the Storage Mirroring connections. You need to configure this resource through the Cluster Administrator.

- Select **Start, Programs, Administrative Tools, Cluster Administrator**.
- Right-click the application group you were working with earlier and select **New, Resource**.
- Specify the following fields on the New Resource dialog box:



- ◆ **Name**—Specify a name that indicates this is the Storage Mirroring virtual server connection.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Double-Take Source Connection**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
- Click **Next** to continue.
 - Verify that both nodes appear as **Possible Owners** and click **Next** to continue.
 - To keep the Double-Take Source Connection resource from coming online before the physical disk, make this resource dependent on the physical disk resource created in **Configure the virtual server on the cluster** on page C-13. Click **Next** to continue.

7. Specify the following on the Double-Take Source Connection Parameters dialog box:

Double-Take Source Connection Parameters

Double-Take Virtual Server Connection

Specify the replication set and Double-Take target to connect.
To setup intra-cluster replication, specify a cluster node as the Double-Take target.

Replication Set (case sensitive): Rep Set Name

Double-Take Target: 172.16.12.47

☐ Allow connection scripts to interact with the desktop

☒ Logon to Target as:

UserName: administrator

Password: XXXXXXXXXX

Domain: domain.com

< Back Finish Cancel

- ◆ **Replication Set**—Specify the name of the Storage Mirroring replication set. This name is case-sensitive and should be the same name as specified in [Configure the Storage Mirroring replication set](#) on page C-17.
- ◆ **Double-Take Target**—Specify the name of the standalone target server.
- ◆ **Allow connection scripts to interact with the desktop**—Mark this check box if you want to display the connection information in a command prompt dialog box.

NOTE: If the script is set to interact with the desktop, the results will be displayed on the owning node. The resource will remain at the **Online pending** status until the script is manually acknowledged by a user. Additionally, if you administer the cluster through Terminal Services, the command prompt windows will not appear on your Terminal Services client, only on the nodes themselves. Therefore, this option should not be a permanent setting, but should be used for troubleshooting if you are having problems establishing a connection.

For additional information on the scripts used, see [Understanding the Double-Take Source Connection resource](#) on page C-27.

- ◆ **Logon to Target as**—Allows you to specify an ID, password and domain to be used when logging on to the target.

NOTE: The user must be a member of the **Double-Take Admin** security group on all nodes of the cluster running a Double-Take Source Connection resource.

If you do not specify a logon, the ID used to start the cluster service will be used. Verify that the ID has **Double-Take Admin** access.

8. Click **Finish** to complete the creation of the Double-Take Source Connection resource.

At this point, you are ready to bring the components online. But you cannot do that until the standalone target is up and running. Therefore, you need to configure the target and then you will come back and bring the components online on the cluster.

Configuring the standalone target

If you have not already done so, install Storage Mirroring on the target using the installation defaults.

NOTE: The Storage Mirroring installation will prompt you to reboot the machine. You can reboot at a later time, but in order for Storage Mirroring to function properly and for you to finish these instructions, you must reboot.

For complete installation details, see the *Getting Started* guide.

The target machine should have drives or volumes with the same drive letters and are of the same approximate size as the drives or volumes that are to be mirrored and replicated from the cluster nodes. No other special configurations is required for the target machine, other than Storage Mirroring be up and running.

Completing the cluster configuration

You will finalize the cluster configuration by bringing the components online.

Bring the Components Online

From the Cluster Administrator (**Start, Programs, Administrative Tools, Cluster Administrator**) bring the Double-Take Source Connection resource and the virtual server resources online.

Notice the warning signs are removed from the resource and the group. The **State** column indicates the resource is online and functioning properly. The data is now being replicated from the source cluster to the target cluster.

NOTE: If the State column indicates failed, as it did prior to bringing the resource online, the connection has not executed properly. When the resource has failed, the Storage Mirroring connection is not active and your data is not being protected in the event of a disaster.

To troubleshoot a failed Source Connection resource, right-click the resource and select **Properties**. Select the **Parameters** tab. Mark the **Allow connection script to interact with desktop** checkbox and click **OK**. When you attempt to bring the resource online, the results from the connection script will be displayed on screen on the node running the script allowing you to see what is causing the failure.

When using the **Allow connection script to interact with desktop** option with the Double-Take Source Connection resource, the resource will report a pending state until the command prompt is closed. The actual status of the resource can be obtained by viewing the script output in the command prompt. Additionally, closing the command prompt by clicking on the close button in the upper right hand corner of the prompt will cause the resource to report a failure. Closing the command prompt by pressing a key on the keyboard ensures that the resource will report its correct state. Additionally, if you administer the cluster through Terminal Services, the command prompt windows will not appear on your Terminal Services client, only on the nodes themselves.

For additional information on the scripts used, see [Understanding the Double-Take Source Connection resource](#) on page C-27.

If you need to terminate the Storage Mirroring connection, right-click the resource and select **Take Offline**. The resource will appear, offline, in the Cluster Administrator. If the Double-Take Source Connection resource is offline, data is not being replicated and is not protected in the event of a disaster.

Standalone to Cluster Configuration

The following instructions are for a standalone source server to a cluster target configuration.

Configuring the standalone sever

1. If you have not already done so, install Storage Mirroring on the standalone server using the installation defaults.

NOTE: The Storage Mirroring installation will prompt you to reboot the machine. You can reboot at a later time, but in order for Storage Mirroring to function properly and for you to finish these instructions, you must reboot.

For complete installation details, see the *Getting Started* guide.

2. Start the Storage Mirroring Management Console by selecting **Start, Programs, Storage Mirroring, Management Console**.
3. Double-click on the standalone server on the left pane of the Management Console to login.
4. Right-click the node and select **New, Replication Set**.
5. Enter a name for the replication set and press **Enter**.
6. On the right pane of the Management Console, select the directory and files that you want to protect.
7. Right-click the replication set name and select **Save**.

Configuring the cluster


For the cluster you will need to configure the virtual server and install Storage Mirroring.

Configure the virtual server on the cluster

Client machines need to access a server when looking for an application, printer, file share, or other resource. Network clients will not be aware which node owns the resources. Therefore, a virtual server needs to be created. This virtual server must have an IP address, a network name, and a disk resource. The virtual server will move between nodes as necessary, transparent to the clients. The clients will always be able to see and access the resource hosted on the virtual server, no matter which node owns it.

To create a virtual server you must create a resource group and add to it an IP address, network name, and a physical disk.

1. Create a resource group for the virtual server.
 - a. Select **Start, Programs, Administrative Tools, Cluster Administrator** and connect to the first cluster.
 - b. Right-click the **Groups** folder on the left pane of the Cluster Administrator and select **New, Group**.
 - c. Specify the **Name** and **Description** and click **Next** to continue.



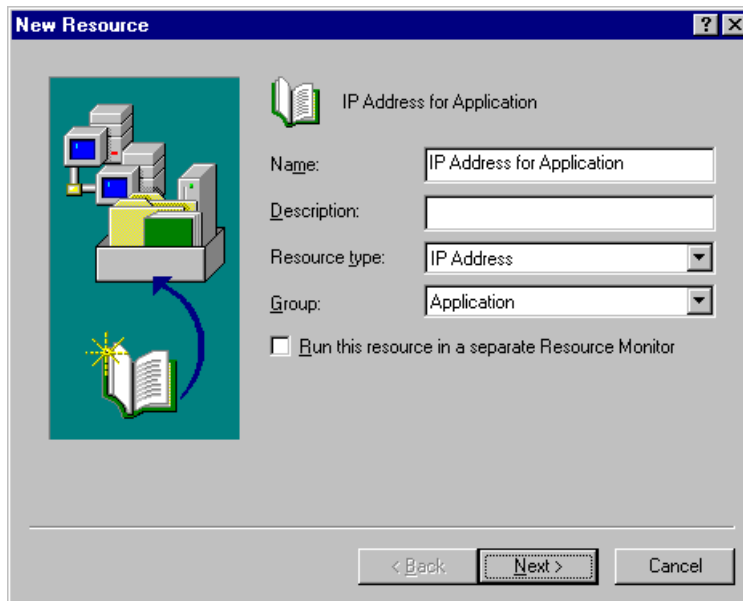
- d. Specify zero or one cluster node as the **Preferred Owner**. Do not specify both nodes. If a preferred node is selected for a resource group, then the group will automatically move to its preferred node when that node comes online. This does not affect manual moves of the group when both nodes are online.

NOTE: For the Preferred Node functionality to work for a resource group, Failback must be also be enabled on the Properties Failback tab of that resource group.

- e. Click **Finish** to complete the creation of the new group.

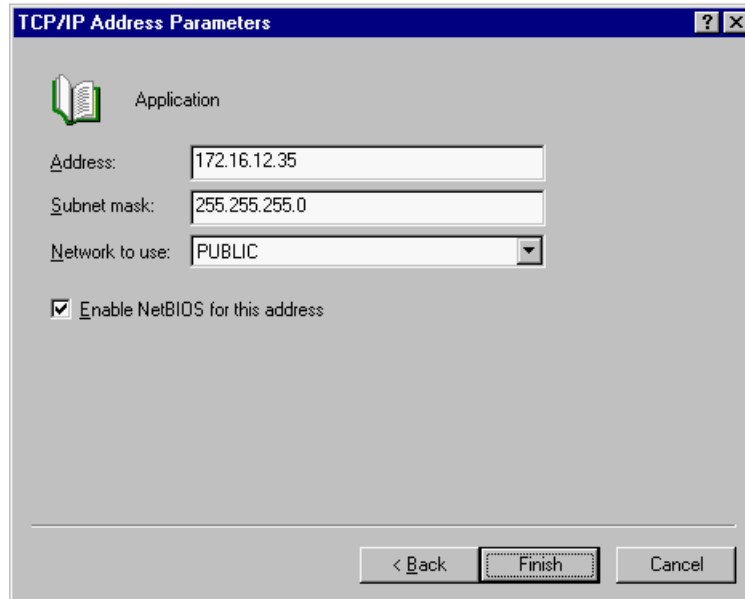
NOTE: You will be notified that the group was created successfully. Click **OK** to acknowledge the message and return to the Cluster Administrator main screen.

- 2. Create an IP Address resource for the new group.
 - a. Right-click the group that you just created and select **New, Resource**.
 - b. Specify the following fields on the New Resource dialog box:



- ◆ **Name**—Specify a name that identifies this resource as the IP address for the group. This name must be unique within the cluster.
- ◆ **Description**—You can optionally add a more detailed description for this resource.
- ◆ **Resource type**—Specify **IP Address**.
- ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
- c. Click **Next** to continue.
- d. Verify that both nodes appear as the **Possible Owners** and click **Next** to continue.
- e. There are no **Dependencies** required for an IP address resource. Click **Next** to continue.

- f. Specify the following fields on the TCP/IP Address Parameters dialog box:

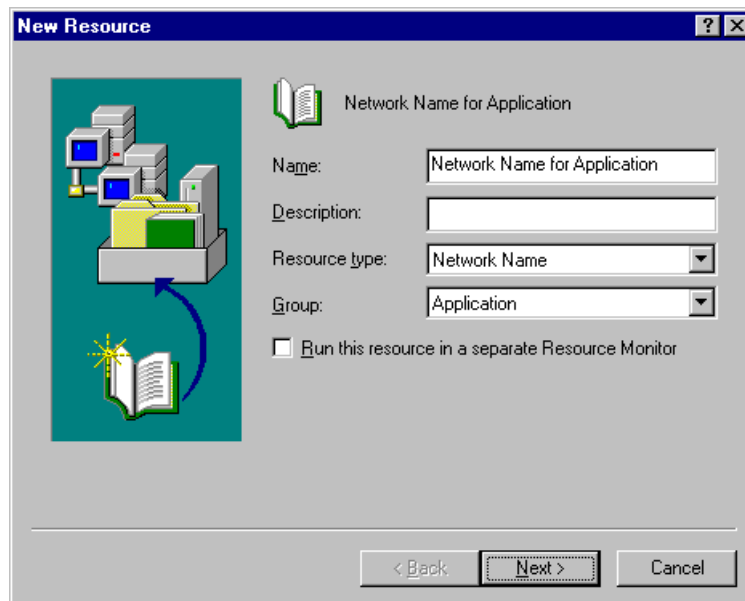


- ◆ **Address**—Enter the static IP address that will be assigned to the application's virtual server.
- ◆ **Subnet mask**—Enter the subnet mask associated with the IP address you just entered.
- ◆ **Network to use**—If you have more than one route for network traffic defined, specify the network that this IP address will use. This should be the same network that the clients are using.

- g. Click **Finish** to complete the creation of the IP Address resource.

3. Create a Network Name resource for the new group.

- a. Right-click the group and select **New, Resource**.
b. Specify the following fields on the New Resource dialog box:

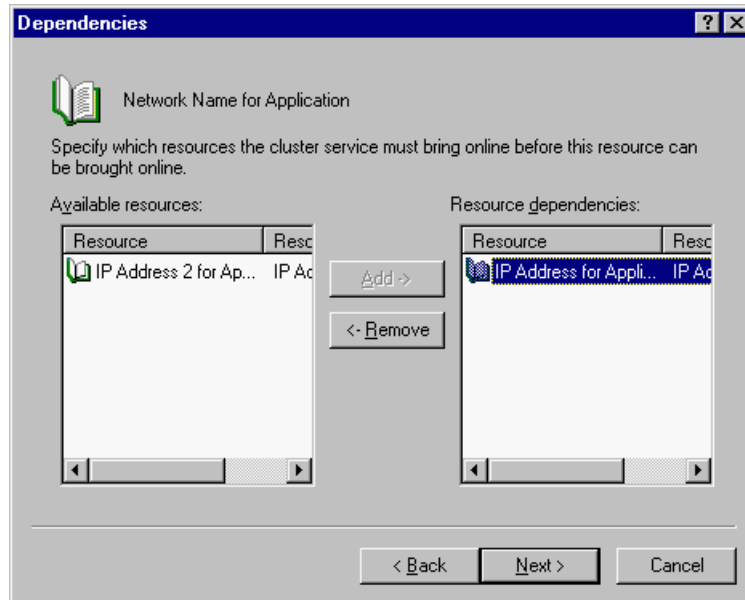


- ◆ **Name**—Specify a name that identifies this resource as the virtual server name. This name must be unique within the cluster.
- ◆ **Description**—You can optionally add a more detailed description for this resource.
- ◆ **Resource type**—Specify **Network Name**.
- ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.

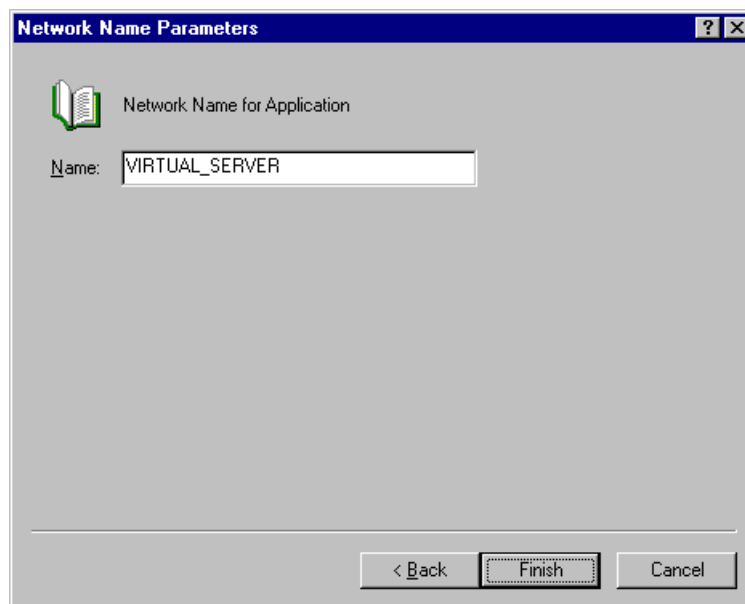
- c. Click **Next** to continue.

- d. Verify that both nodes appear as **Possible Owners** and click **Next** to continue.

- e. An IP address must be present in order for a network name to be assigned. Therefore, move the IP address resource associated with the virtual server to the **Resource dependencies** list and click **Next** to continue.




- f. Specify the Network Name Parameters by entering the virtual name of the server. This is the name that clients will look for on the network and the name that will register with DNS. It must be unique both on the cluster and on the network.



- g. Click **Finish** to complete the creation of the Network Name resource.
4. Create a Physical Disk resource for the new group.
- Right-click the group and select **New, Resource**.
 - Specify the following fields on the New Resource dialog box:
 - ◆ **Name**—Specify a name that identifies the disk drive associated with the virtual server. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Physical Disk**.
 - ◆ **Group**—The resource group name should be selected. If it is not, select the correct group name.
 - Click **Next** to continue.
 - Verify that both nodes appear as **Possible Owners** and click **Next** to continue.

-
- e. No other resources are required as dependencies for disk resources. Click **Next** to continue.
 - f. Specify the disk drive associated with the physical disk resource.



- g. Click **Finish** to complete the creation of the Physical Disk resource.

The application group now contains all of the resources necessary to identify a virtual server. All resources should be offline at this time. If any are online, they need to be taken offline by right-clicking the resource and selecting **Take Offline**.

Install Storage Mirroring on the cluster

1. If you have not already done so, install Storage Mirroring on the first node of the cluster using the installation defaults.

NOTE: The Storage Mirroring installation will prompt you to reboot the machine. You can reboot at a later time, but in order for Storage Mirroring to function properly and for you to finish these instructions, you must reboot.

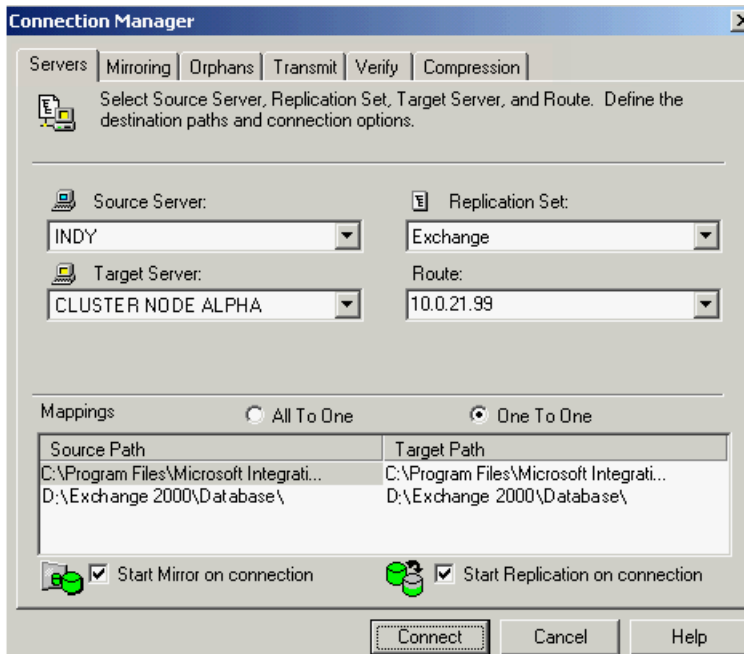
For complete installation details, see the *Getting Started* guide.

2. Repeat step 1 on the second node of the cluster.

No special configurations is required for the cluster nodes, other than Storage Mirroring be up and running.

Completing the standalone configuration

To finalize the configuration, you need to establish a connection between the replication set on the source and the cluster. Using the Connection Manager, establish the connection.



Select the owning node as the **Target Server** and select the IP address of the virtual server for the **Route**. For detailed information on establishing a connection through the Connection Manager, see *Establishing a connection on page 6-7*.

Understanding the Double-Take Source Connection resource

The Double-Take Source Connection resource is like other cluster resources. The cluster controls the movement from node to node if there is a failure. This resource additionally controls the Storage Mirroring connection. Therefore, if there are node failures, the Storage Mirroring connection is automatically handled by the new owning node without user intervention.

When you are first setting up your Storage Mirroring configuration or when troubleshooting, you may want to understand what the resource is doing. When the Double-Take Source Connection resource is brought online or taken offline a command shell environment is created and a Storage Mirroring DTCL script is executed. The script either establishes a Storage Mirroring connection or disconnects the Storage Mirroring connection, depending on the node. If you are using the Double-Take Source Connection resource, the cluster object will control these scripts as the cluster service switches between nodes. If you do not use the Double-Take Source Connection resource, there is no mechanism to bring a Storage Mirroring connection online. The non-owning node does not know that a connection was present because the connection information stored in the connect.sts is not replicated between systems. Therefore, if that non-owning node becomes the owner, the connection would need to be manually re-established. So it is important to always use the Double-Take Source Connection resource when the source is a cluster configuration.

Outlined briefly below are the steps used in the DTCL online and offline scripts.

Online Script

1. Identify the elements used in the script: source, target, replication set, and so on.
2. Login to the target.
3. Specify the source.
4. Establish a connection from the source to the target using identical drive mappings on the target as on the source.

NOTE: Notice that identical drive mappings are used in the online script. If your source and target do not have the same drive letters, the online script will fail.

5. Begin a block checksum mirror.

The source will replicate changes to the data after the mirror completes.

Offline Script

1. Identify the elements used in the script: source, target, replication set, and so on.
2. Login to the target.
3. Specify the source.
4. Determine the connection ID and disconnect it.

If you want to view the DTCL script files directly, they are called `online.dtcl` and `offline.dtcl` and are located in the directory where you installed Storage Mirroring. These are text files that can be viewed with any text editor.

NOTE: Making changes to the `online.dtcl` or `offline.dtcl` scripts is not recommended. Those users who are experienced and comfortable creating and editing Storage Mirroring scripts may want to modify them on their own. Such modifications may not be supported by technical support.

Storage Mirroring can be used in a Network Address Translation (NAT) environment provided that the framework is established properly. This appendix describes NAT basics and provides instructions for configuring your Storage Mirroring servers for a NAT environment. You must have a router in place and know how to configure the ports. If you do not have a NAT router installed or if you do not know how to configure the ports, see your NAT reference manual.

NAT Overview

Network Address Translation (NAT) enables a network to pool a set of internal IP addresses while presenting them as one address externally to the Internet. For example, NAT provides a company the ability to use more internal IP addresses. Since they are only used internally, there is no possibility of conflicts with IP addresses used by other companies and organizations. NAT might also provide a firewall by hiding internal IP addresses. The functionality provided by NAT depends on your configuration.

- ◆ **Static NAT**—This is the most simple configuration where a single, internal IP address is always mapped in a one-to-one correlation to a single, external IP address. If your NAT environment only has two Storage Mirroring servers, each on either side of a router, you need to establish static mapping on the router so that the two Storage Mirroring machines can communicate.
- ◆ **Dynamic NAT**—This configuration is dynamic because it allows a single, internal IP address to be mapped to any one of a group of external IP addresses managed by the router. The router maintains the external IP address selected and assigned from the group so that any incoming packets destined for that IP address are routed to the correct internal IP address. Storage Mirroring is not suited for this environment because of its automated features, such as auto-reconnect that attempts to automatically reestablish connections after a failure. With dynamic IP addresses, these features will not work.
- ◆ **Overloading NAT**—This configuration is known by various names including NAPT (network address port translation) and PAT (port address translation). It allows multiple, internal IP addresses to be mapped to a single, external IP address. The external IP address is distinguished by an appended port number. That makes the internal IP address to external IP address: port number mapping unique. If your NAT environment has multiple Storage Mirroring machines on one side of a router, you will have to establish port mappings for both the router and for Storage Mirroring.

NOTE: If you are using overlapping NAT, where the same IP address is being used on both sides of the router, you can use a static or overloading NAT configuration, but do not use dynamic.

Storage Mirroring and NAT Overview

Regardless of whether you are using a static or overloading NAT environment, you need to configure your environment so that Storage Mirroring traffic is permitted access through the router and directed appropriately. If you are using a static NAT environment, your setup might look similar to figure 1.




Figure 1: Sample static NAT scenario with one Storage Mirroring server on each side of a router.

If you are using an overloading NAT environment, your setup might look similar to figure 2.




Figure 2: Sample overloading NAT environment with unique ports.

Modifying the Storage Mirroring Ports

Storage Mirroring uses multiple ports for communications. By default, during the installation, Storage Mirroring automatically sets these ports to an optimal configuration for a non-NAT environment. You will need to modify your port settings for NAT integration. The following table outlines all of the Storage Mirroring ports and gives brief guidelines on when to modify each port. Make the port modifications necessary according to your specific environment.

Port Name	Management Console Name	Communication Type	Default Port Number	When to Modify the Port
DirUNetPort NetPort	Status Listen Service Listen	Directed UDP TCP	1105 1100	<ul style="list-style-type: none">♦ Static NAT—If you are using a static NAT environment where you only have one Storage Mirroring server on each side of the router, you do not need to modify your Storage Mirroring ports. You can keep the default port selection as it is for both Storage Mirroring machines. You will need to add a mapping to your router for the Storage Mirroring port.♦ Overloading NAT—If you are using an overloading NAT environment where you have more than one Storage Mirroring server on at least one side of the router, you must set unique Storage Mirroring ports. You may keep the default port selection for one server, but subsequent servers must have different port numbers. From the Management Console, login into a server and access the Server Properties, Network tab (right-click the server and select Properties). Modify the Service Listen and Status Listen ports so that each Storage Mirroring server uses unique ports. These two ports need to have the same number on a single server. The port numbers must be unique from server to server. Each port number that you use will need to be added to the router mapping.
UNetPort	Heartbeat Transmit	UDP	1100	This port is used to control the auto-population of the Management Console server tree. It does not have an impact on a NAT environment since servers across a router have to be manually inserted. Therefore, there is no need to modify this port.
StatsPort		TCP	1106	<ul style="list-style-type: none">♦ If you want to run one instance of DTStats across a router, you will need to add a mapping to your router for the StatsPort.♦ If you want to run DTStats for multiple machines, each Storage Mirroring server must have a unique StatsPort and each port will need to be added to the router mapping.♦ To modify the StatsPort, use the <code>DTCL set</code> command with the <code>statsport</code> option.♦ In order for DTStats to function properly, the StatsPort cannot be set to the same value as the DirUNetPort / NetPort.

NOTE: You must stop and restart the Storage Mirroring service for a changed port setting to take effect.

Record your Storage Mirroring port settings and the corresponding IP address of that server.

Storage Mirroring Port Configuration

Server Name	Port	IP Address

Configuring the Router

Regardless of whether you are using a static or overloading NAT environment, you need to configure your router so that Storage Mirroring traffic is permitted access through the router and directed appropriately. Using the information you recorded in the port configuration table above, set up the configuration table on the router to identify each server and its corresponding IP address, Storage Mirroring port and NAT port.

NOTE: Make sure that both the UDP and TCP ports on the router match the port used by the Storage Mirroring server.

Since Storage Mirroring communication occurs bidirectionally between the source and target servers, make sure you configure your router(s) for all of your Storage Mirroring machines.

Storage Mirroring failover use ICMP pings or UDP heartbeats to determine if the source server is online. If you are using ICMP pings and a network device between the source and target is blocking ICMP traffic, failover monitors cannot be created or used. In this situation, you must configure your network devices to allow ICMP pings between the source and target.

Since there are many NAT routers on the market, each router can be configured differently. See your NAT reference manual for instructions on setting up your particular router.

Server	Public Address	Private Address	UDP and TCP Router Ports	HPOVSM Port
Jersey	10.0.0.2	192.168.1.3	1205	1205
Philly	10.0.0.2	192.168.1.4	1305	1305

Figure 3: Sample router table.

Using the Management Console in a NAT Environment

By default, servers across a router will not automatically be populated in the Management Console. You must insert them manually. After the server is inserted, connections are established like other Storage Mirroring connections but you need to specify the correct route. The next two sections cover these topics in detail.

Inserting the servers in the Management Console

When you start the Management Console, it will not be able to see servers automatically across a router. Therefore, those machines will need to be manually inserted.

NOTE: If a machine is manually inserted into the Management Console, it will automatically be saved in your workspace and will appear the next time the Management Console is started. If you established unique Storage Mirroring ports, subsequent restarts of the Management Console will display the server saved in the workspace, but it will appear with a red X as if there is a problem. The X appears because communication cannot automatically be established. If you login to the server (by double-clicking on it), communications will be established and the X will be removed.

1. Open the Management Console (**Start, Programs, Storage Mirroring, Management Console**).
2. Select **Insert, Server**.
3. Type the IP address of the router and the port number used by the Storage Mirroring server.

NOTE: If your traffic is routed across more than one router, enter the IP address of the router that the Storage Mirroring server is connected to. For example, if you are using the Management Console as displayed in Figure 2 and need to insert the target Indy, you would insert the IP address 10.0.0.1 and the port number 1105.

4. Click **Test**. After the machine is located, the machine name will be inserted into the Management Console server tree.

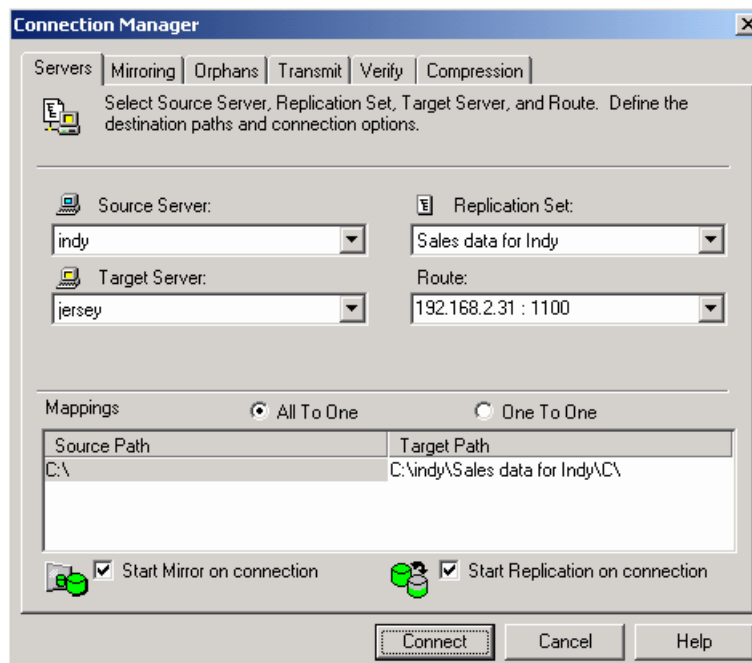
WARNING: If the machine was located, but Storage Mirroring is not running on it, you will be given the choice of inserting the machine anyway. Click **Yes** to insert the server into the Management Console server tree. Verify that Storage Mirroring is running on that machine before continuing.

If the server could not be located, you will see a message to the right of the **Test** button stating that it was not found. Verify that you are entering the correct IP address and that you have configured your router correctly.

Establishing the connection

Connections across a router are established like any other Storage Mirroring connection, but you must make sure that you specify the correct route across the network so that the Storage Mirroring traffic will be transmitted successfully.

1. Open the Connection Manager.
2. Some entries on the **Servers** tab will be completed depending on which method you used to access it. Verify that the correct **Source Server** and **Replication Set** are identified
3. Select your **Target Server** and the IP address of that machine will automatically be inserted in the **Route** field.
4. Since the router is going to prevent sending data directly to the IP address of the target machine, change the **Route** to the IP address and port of the router.



NOTE: If your traffic is routed across multiple routers, enter the IP address of the router that the Storage Mirroring server is connected to. For example, if you are using the Management Console as displayed in Figure 2 and need to connect to the target Indy, you would select 10.0.0.1:1105.

If your target machine is located on the same subnet as your Management Console, the route field will not automatically contain the target machine's router information. For example, if the Management Console in Figure 2 was on the target subnet instead of the source subnet, you would have to manually type in 10.0.0.1:1105 in the **Route** field rather than selecting it from the list.

5. Select the target path, which corresponds to the desired location on the target where the replicated data will reside, by selecting **One-to-One** or **All-to-One**. If you want to specify a custom location, select either of the radio buttons and then click on the directory entry under the **Target Path** column. You will be in edit mode and can specify the desired location on the target for the replicated data.
6. If you want mirroring and/or replication to start immediately when the connection is established, mark either or both of the two check boxes at the bottom of the **Servers** tab.

NOTE: Other tabs are available in the Connection Manager to set advanced connection settings. To establish a connection, you do not need to make modifications to these tabs; a connection can be established with the default settings.

7. Click **Connect** to establish the connection.

Your connection is now established. Data will be transmitted to the target across the router. Storage Mirroring functionality and monitoring will continue and be available as with any other type of Storage Mirroring connection.

Using the Text Client in a NAT Environment

Since servers are not automatically populated in the Text Client, you do not have to worry about inserting them as you did in the Management Console. But you still need to login to the machine and specify the correct route to establish a connection.

1. Login to the source by using the `login` command. This command is the same except that you need to specify the IP address and port of the machine, rather than the machine name.

Command	<code>LOGIN</code>
Description	Log on to a Storage Mirroring machine
Syntax	<code>LOGIN <machine> <username> <password> [domain]</code>
Options	<ul style="list-style-type: none">◆ machine—Name of the machine◆ username—Name of the user◆ password—Password associated with username.◆ domain—If logging in using a domain account, this is the domain name. If logging in using a local account, this is the machine name.
Examples	<code>login indy administrator *****</code>
Notes	<ul style="list-style-type: none">◆ The <code>login</code> command is not available when scrolling through the Text Client command history.◆ If characters in the password include non-alphanumeric characters, the password field must be enclosed in quotation marks.◆ The password cannot be a Storage Mirroring keyword. These are any DTCL command (source, target, and so on.) or any DTCL shortcut command (env, mon, rep, and so on).

2. Identify the machine that you just logged into as the source by using the `source` command. This command is also the same except that you need to specify the IP address and port of the machine, rather than the machine name.

Command	<code>SOURCE</code>
Description	Identifies a machine as the active source machine
Syntax	<code>SOURCE <source_machine></code>
Options	source_machine —Name of the machine
Examples	<code>source indy</code>

3. Connect the replication set to the target by using the `connect` command. Again, this command is the same except that you need to specify the IP address and port of the target machine.

Command	<code>CONNECT</code>
Description	Establishes a connection between a replication set and a target machine
Syntax	<pre>CONNECT <repset> TO <target_machine> MAP EXACT MAP BASE <target_path> MAP <source_path> TO <target_path> [,...] [MIRROR NOMIRROR] [, REPLICATE NOREPLICATE] [, MONITOR NOMONITOR] [, ORPHANS NOORPHANS] [, COMPRESSION <level>] [CLEARRESTOREREQUIRED]</pre>

Options

- ◆ **repset**—Name of the replication set
- ◆ **target_machine**—Name of the target machine, an IP address on the target machine, or a virtual IP address
- ◆ **MAP EXACT**—Specifies that the replication set data will be sent to the same logical volume on the target (c:\data and d:\files is copied to c:\data and d:\files, respectively)
- ◆ **MAP BASE target_path**—Substitute a complete path, including the volume, for **target_path** and the data will be replicated to target_path\SrcVolName on the target machine
- ◆ **MAP source_path TO target_path**—Custom location that specifies each directory on the source and where that data will be copied to on the target machine
- ◆ **...**—Indicates that the **source_path TO target_path** option can be used more than once for each source directory in the replication set
- ◆ **MIRROR**—Automatically initiates a mirror when the connection is established
- ◆ **NOMIRROR**—Does not initiate a mirror when the connection is established
- ◆ **REPLICATE**—Automatically initiates replication when the connection is established
- ◆ **NOREPLICATE**—Does not initiate replication when the connection is established
- ◆ **MONITOR**—Specifies that the target is going to monitor the specified source machine for failover. The source machine must have already been defined as a monitor machine.
- ◆ **NOMONITOR**—Specifies that the target is not going to monitor the source machine for failover
- ◆ **ORPHANS**—Removes orphan files on the target
- ◆ **NOORPHANS**—Does not remove orphan files on the target
- ◆ **COMPRESSION level**—Sets the compression level for data being sent to the target using levels 1 through 3 with 1 being minimum compression and 3 being maximum compression. This option sets the level only. It does not enable compression.
- ◆ **CLEARRESTOREREQUIRED**—Clears the restore required flag and initiates the connection

Examples

- ◆ connect Exchange to jersey map exact
- ◆ connect sql to jersey map base d:\DTFiles\

Notes

- ◆ The default settings for this command are mirror, replicate, nomonitor, and noorphans.
- ◆ If a path begins with a non-alphabetic character or if the replication set name includes non-alphanumeric characters, the path or name must be enclosed in quotation marks.
- ◆ If you are establishing a connection within a NAT environment, you will need to specify the port of the router after the IP address (seperated by a colon). See [NAT Environments](#) on page D-1 for more details.

Your connection is now established. Data will be transmitted to the target across the router. Storage Mirroring functionality and monitoring will continue and be available as with any other type of Storage Mirroring connection.

A firewall is a program that helps prevent unauthorized access to or from a private network. A firewall examines each packet entering or leaving the network and blocks those that do not meet the specified security criteria.

Storage Mirroring can be used through a firewall as long as the firewall ports are configured to allow Storage Mirroring servers to communicate with each other. This appendix describes the ports used for Storage Mirroring communications, how to verify the port settings for your Storage Mirroring servers, and the proper firewall port configuration for use with Storage Mirroring. This section assumes that you have a firewall program installed and you know how to open your firewall ports.

Verifying Storage Mirroring Port Settings

Storage Mirroring uses specific ports for communication between the source, target, and client machines. In order to use Storage Mirroring through a firewall, you must first verify the current Storage Mirroring port settings so that you can open the correct ports on your firewall to allow Storage Mirroring machines to communicate with each other.

Using the following table, locate and record your port settings for each of the four Storage Mirroring ports. As the table indicates, there are three different Storage Mirroring clients (Management Console, Text Client, and Failover Control Center) that you can use to verify each of the four port settings. For more specific information on using a particular Storage Mirroring client to verify a port setting, see the chapter named for that client in this manual.

Port Name	Additional Names	Communication Type and Port Usage	Use Any of These Clients to Verify the Storage Mirroring Port Setting	Default Port Setting	Record Your Port Setting
NetPort	<ul style="list-style-type: none"> ◆ Service Listen Port ◆ Service Transmit Port 	The NetPort is used for TCP communication between Storage Mirroring servers and clients.	<ul style="list-style-type: none"> ◆ Management Console— See the Network tab on the Server Properties dialog box. ◆ Text Client—Use the <code>get netport</code> command. ◆ Failover Control Center—See the Set Communication Parameters dialog box. 	1100	
UNetPort	<ul style="list-style-type: none"> ◆ Heartbeat Advertisement Port ◆ Heartbeat Transmit Port ◆ Heartbeat Listen Port 	The UNetPort is used to send and receive Storage Mirroring heartbeats which are broadcast UDP communications.	<ul style="list-style-type: none"> ◆ Management Console— See the Configuration tab on the Management Console Options dialog box. ◆ Management Console— See the Network tab on the Server Properties dialog box. ◆ Text Client—Use the <code>get unetport</code> command. ◆ Failover Control Center—See the Set Communication Parameters dialog box. 	1100	
DirUNetPort	<ul style="list-style-type: none"> ◆ Status Transmit Port ◆ Status Listen Port 	The DirUNetPort sends directed UDP communications to request and receive status updates, keeping the Storage Mirroring Management Console at-a-glance monitoring current.	<ul style="list-style-type: none"> ◆ Management Console— See the Configuration tab on the Management Console Options dialog box. ◆ Management Console— See the Network tab on the Server Properties dialog box. ◆ Text Client—Use the <code>get dirunetport</code> command. 	1105	
StatsPort		The StatsPort accepts TCP communications from DTStat.	<ul style="list-style-type: none"> ◆ Text Client—Use the <code>get statsport</code> command. 	1106	

NOTE: Because these ports appear on multiple machines (sources, targets, and clients), the port settings have to be identical on the machines that need to communicate with each other. For example, a source with a NetPort setting of 1100 will only be able to communicate with other machines (sources, target, and clients) that also have a NetPort of 1100. If any machine has a NetPort setting of any value other than 1100, that machine cannot communicate with this source. Therefore, if you change a port setting on one machine, perhaps your source, you will also need to make the same change to the target(s) and client(s) that the source needs to communicate with.

Opening the Firewall Ports

Using the values from the previous table, configure your firewall ports for both inbound and outbound traffic so that the Storage Mirroring servers can communicate with each other. For specific steps in configuring these ports, see your firewall reference manual.

1. Open the port assigned to the NetPort and StatsPort for the inbound and outbound transmission of TCP packets.
2. Open the port assigned to the UNetPort and DirUNetPort for the inbound and outbound transmission of UDP packets.

NOTE: Storage Mirroring failover use ICMP pings or UDP heartbeats to determine if the source server is online. If you are using ICMP pings and a network device, such as a firewall, between the source and target is blocking ICMP traffic, failover monitors cannot be created or used. In this situation, you must configure your firewall to allow ICMP pings between the source and target.

Alerts—Notifications, warnings and errors that describe Storage Mirroring processes.

Auto-disconnect—When the Storage Mirroring queue has reached its maximum, the auto-disconnect process empties the source queue and the source/target connection is broken.

Auto-reconnect—The process that automatically reestablishes any source/target connections that were active before a Storage Mirroring failure on the source machine.

Auto-remirror—If a source or target machine fails, auto-remirror automatically initiates the mirroring process after the target machine is back online or an auto-reconnect has occurred.

Block Checksum Comparison—Formula applied to two blocks of data to determine if the binary make-up of the blocks are identical. Also called a checksum comparison.

Cluster—A group of independent servers, called nodes, working together as a unit.

Connection—The link between a source replication set and a target machine. This is a logical connection that refers to the relationship between the replication set on the source and the copy of the data on the target, as well as the status of the connection.

Connection ID—A unique ID associated with each Storage Mirroring connection.

Drive Shares—Any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users.

Failover Timeout—The amount of time before failover begins. The time is calculated by multiplying the values of the monitor interval and missed packets settings, and is displayed in the Failover Control Center.

File Change—The creation, modification, or deletion of a file. Also called a file request or operation.

Groups—Groups exist on a cluster to provide an administrative method for collecting multiple resources together for simple management.

LogViewer—The Storage Mirroring logging utility that logs processing notifications, warnings, and errors. The logs are viewed with the LogViewer.

MAC Addresses—Media Access Control addresses represent the physical station or hardware address of a NIC. These addresses are used to identify computers on a network. MAC addresses are unique and are coded into every network adapter by the manufacturer.

Mirror—Mirroring is the process of transmitting data contained in a replication set from the source to the target machine so that an identical copy of data exists on the target machine.

Missed Packets—The configurable setting that specifies how many monitor replies can be missed before assuming a source machine has failed.

Monitored IP Address—Those IP addresses on a source that have been selected to be monitored for failure by a Storage Mirroring target.

Monitor Interval—The configurable failover setting that specifies how often the monitor request is sent to a source machine.

Monitor Reply—The acknowledgement that a source sends to a target indicating that it is still online.

Monitor Request—An inquiry sent from a target to a source to determine if the source is still online.

MSCS—Microsoft Cluster Service. This is the Microsoft clustering solution. A cluster is a group of independent servers, called nodes, working together as a unit.

Node—An independent server that is one of the units that make up a cluster.

One-to-One Configuration—Configuration used when one target machine is dedicated to support one source machine for mirror and replication data.

One-to-One Connection—Mirror and replication data are copied to the same volume and directory structure on a target as they exist on a source. For example, c:\data and d:\files on the source are copied to c:\data and d:\files, respectively, on the target.

Queue—A first-in, first-out pool of operations on the source waiting to be transmitted to a target, or a first-in, first-out pool of operations waiting to be written to disk on the target.

Quorum—The quorum resource is used on a cluster to guarantee that only one active node operates as the cluster, acting as a tie-breaker in the unlikely event that cluster nodes are running but cannot communicate.

Recursion—The process of automatically applying a replication set rule to the subdirectories of the named directory.

Remirror—Repeating a mirror process to guarantee the integrity of the data on the target.

Replication—The real-time transmission of file changes.

Replication Set—The data on a source machine that is to be protected by Storage Mirroring. Replication sets are defined by the volumes, directories, files, or wild card combinations that are to be replicated to the target.

Replication Set Rule—One volume, directory, file, or wild card specification that is to be replicated to the target. One or more replication set rules form a replication set.

Resource—A physical or logical unit that is managed by a cluster and allowed to move from one machine to the other. The shared SCSI disk used in a MSCS cluster is considered a resource as well as drive shares, services, applications, IP addresses, computer names, and so on.

Share—Any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users.

Source—The machine that contains the original mission-critical data that must be protected.

Synchronize—The mirror process that copies the replication set data from the source to the target.

Target—The machine that maintains the copy of the replication set from the source.

Unmonitored IP Addresses—On a machine with multiple IP addresses, Storage Mirroring will not monitor for failure the IP addresses that are not selected, but the target can assume these addresses during failover, if configured.

Verification—The process of confirming that the Storage Mirroring data on the target is identical to the replication set data on the source.

Index

A

- Active Directory** 15-7
- active/active** 1-8
- active/standby** 1-7
- administrator access** 15-1
 - Management Console 2-6
 - text clients 3-4
- advanced settings**
 - Management Console
 - database 14-10
 - general 14-2
 - logging 14-11
 - network 14-5
 - notification 14-12
 - queuing 14-6
 - source 14-7
 - startup 14-4
 - target 14-9
 - settings table for text clients 14-15, 14-37
 - text clients 14-14
- all-to-one mapping** 6-2
- application failover**
 - Windows B-2
- at-a-glance monitoring** 13-6
- auto-disconnect** F-1
 - overview 6-7
- AutoReconnect** 6-24
- auto-reconnect** F-1
 - Management Console 6-15, 14-4
 - text clients 6-24
- AutoRemirror** 7-18
- auto-remirror** 6-15, 6-24, F-1
 - Management Console 7-6, 7-10, 14-4
 - overview 7-4
 - text clients 7-18

B

- backup features** 1-2
- bandwidth limitations** 10-2
 - Management Console 10-6
 - text clients 10-12
- block checksum** 7-1, 7-13, F-1
- byte limit** 10-3

C

- calculate replication set size**
 - Management Console 5-10, 7-6, 7-7
 - text clients 5-14
- CalculateOnConnect** 5-14
- chained configuration** 1-11
- checksum** 7-1, 7-13, F-1
- circular configuration** 6-2
- clean** 7-13
- cluster**
 - cluster configurations C-2
 - cluster to cluster configuration C-5
 - cluster to standalone configuration C-13
 - MSCS and Storage Mirroring C-3
 - overview C-1
 - providing disaster recovery C-4
 - resource C-27
 - resources 6-4, F-2
 - standalone to cluster configuration C-21
- Command Line Client**
 - file entry 3-3
 - interactive entry 3-3
 - overview 3-3
 - starting 3-3
 - using 3-3
- commands** A-2
 - conventions A-1
- communication ports**
 - Management Console 2-8
 - text clients 3-7
- compression**
 - Management Console 6-12
 - text clients 6-22
- configurations**
 - chained 1-11
 - many-to-one 1-9
 - one-to-many 1-10
 - one-to-one
 - active/active 1-8
 - active/standby 1-7
 - overview 1-7
 - single machine 1-12
- conid** 6-21, A-3
- connect** 6-16, 7-11, 8-8, A-4
- connect tdu** 6-18, A-5
- connection**
 - auto-disconnect 6-7
 - auto-reconnect
 - Management Console 6-15
 - text clients 6-24
 - cluster resources 6-4
 - disconnecting
 - Management Console 6-9
 - text clients 6-19
 - display
 - Management Console 6-11
 - how it works 6-5
 - ID 6-4, 6-21
 - Management Console
 - auto-reconnect 6-15
 - disconnecting 6-9
 - display 6-11
 - mapping 6-9
 - monitoring 6-10
 - new 6-8
 - queuing 6-13
 - statistics 6-10
 - mapping 6-2
 - Management Console 6-9
 - text clients 6-16
 - monitoring
 - Management Console 6-10
 - new
 - Management Console 6-8
 - text clients 6-16
 - options 6-2
 - overview 6-1
 - queuing
 - Management Console 6-13
 - text clients 6-23
 - route 6-2, 6-8
 - statistics
 - Management Console 6-10
 - text clients
 - auto-reconnect 6-24
 - disconnecting 6-19
 - mapping 6-16
 - new 6-16
 - queuing 6-23
 - wizard 6-8
- connection ID** 6-4, 6-21, F-1
- Connection Wizard** 6-8
- connections**
 - ID F-1
- copying replication set** 5-11
- creating replication set**
 - text clients 5-12

- current session**
 - source/target functionality 3-6
- custom mapping** 6-2, 6-16
- cycle** 6-2

D

- deleting replication set**
 - Management Console 5-11
 - text clients 5-19
- direct rule entry** 5-9
- disaster recovery** 1-2
- disconnect** 6-19, A-5
- disconnecting connection**
 - Management Console 6-9
 - text clients 6-19
- display**
 - Management Console 6-11
 - Text Client 3-2
- displaying replication set** 5-16
- DT_Full_Access** 3-4
- DT_Monitor_Only_Access** 3-4
- DT_No_Access** 3-4
- DTCL commands** A-2
 - conventions A-1
 - scripting A-45
 - examples A-51–A-52
 - flow control A-46
 - variables A-45
- DTCL scripting** A-45
 - examples A-51–A-52
 - flow control A-46
- DTCL variables** A-45
- DTStat** 13-30, 13-34, A-6
 - command prompt 13-34
 - log 14-11
 - Management Console 13-31, 14-4
 - Text Client 13-32
- dynamic volume** 5-6, 6-2

E

- environment** A-10
- Event Viewer** 13-14
- exclude** 5-2, 5-9, 5-11, 5-13, 5-17
- exit** A-10

F

- failback** A-11
 - Failover Control Center 11-24
 - how it works 11-8
 - process 11-3
 - text clients 11-38
- failover** A-11
 - application
 - Windows B-2
 - clients 11-4
 - configuring
 - Failover Control Center 11-17–11-21
 - text clients 11-28–11-36
 - failure detection 11-9
 - hardware configurations 11-12
 - how it works 11-6
 - IP address placeholders 11-25
 - manual intervention 11-27
 - manual share updates 11-26
 - monitoring 4-4, 11-21
 - operating system differences 11-10
 - options 11-15
 - overview 1-5, 11-1
 - process 11-2

- scripts
 - Windows B-2
- terminology 11-5
- testing 11-27, 11-37
- Failover Control Center**
 - configuring failover 11-17–11-21
 - display 4-3
 - display refresh rate 11-26
 - failback 11-24
 - operations 4-3
 - overview 4-1
 - ports 4-6, 11-26
 - starting 4-2
- failure detection** 11-9
- failure monitoring**
 - overview 1-5
- features**
 - backup 1-2
 - disaster recovery 1-2
 - high availability 1-2
 - other 1-2
 - resource utilization 1-2
- file entry** 3-3

G

- get** A-11
- getenvstr** A-12
- getlocal** A-12

H

- heartbeats** 14-5
- help** A-12
 - Management Console 2-12
- high availability** 1-2
- HP OpenView Storage Mirroring**
 - features
 - backup 1-2
 - disaster recovery 1-2
 - other 1-2
 - resource utilization 1-2
 - overview 1-1
- HP OpenView Storage Mirroring**
 - features
 - high availability 1-2

I

- icons**
 - at-a-glance monitoring 13-6
 - connection display 6-11
 - connection statistics 6-11
 - Management Console 2-3, 13-6
- ignore delete operations** 14-7
- include** 5-2, 5-9, 5-11, 5-13, 5-17
- interactive entry** 3-3
- IP address placeholders** 11-25

L

- limit bandwidth** 10-12, A-13
- listing replication set** 5-16, 5-19
- load source** A-13
- load target** A-13
- log**
 - DStat 14-11
 - LogViewer 13-40
 - verification
 - Management Console 9-9, 14-11
 - overview 9-4
 - text clients 9-13
- LogDir** 13-42
- login** 9-11, 12-7, A-14
- logout** A-14
- LogViewer** 13-39, 13-43, A-15

- alert descriptions 13-44
- log file 13-40
- monitoring tools
 - viewing log files F-1
- text clients 13-42

M

Management Console

- advanced settings
 - database 14-10
 - general 14-2
 - logging 14-11
 - network 14-5
 - notification 14-12
 - queuing 14-6
 - source 14-7
 - startup 14-4
 - target 14-9
- at-a-glance monitoring 13-6
- auto-reconnect 6-15
- auto-remirror 7-6, 7-10
- bandwidth limitations 10-6
- calculate replication set size 5-10, 7-6, 7-7
- clear transmission settings 10-5
- compression 6-12
- connection statistics 6-10
- copying replication set 5-11
- deleting replication set 5-11
- disconnecting connection 6-9
- display 6-11
- display colors 13-9
- DStat 13-31
 - log 14-11
- help 2-12
- icons 2-3, 13-6
- LogViewer 13-41
- manually initiating transmission 10-10
- manually initiating verification 9-8
- manually starting replication 8-4
- manually stopping replication 8-4
- mapping 6-9
- message window 13-10
- mirroring on connection 7-6
- monitoring connection 6-10
- monitoring mirroring 7-8
- monitoring replication 8-5
- monitoring tools
 - at-a-glance monitoring 13-6
 - display colors 13-9
 - message window 13-10
 - overview 13-2
 - statistics 13-2
- monitoring verification 9-10
- new connection 6-8
- new replication set 5-7
- orphan files 7-9
- overview 2-1
- pausing mirroring manually 7-7
- permissions 14-7
- ports 2-8
- program settings
 - database 14-10
 - general 14-2
 - logging 14-11
 - network 14-5
 - notification 14-12
 - queuing 14-6
 - source 14-7
 - startup 14-4
 - target 14-9
- queuing 6-13
- renaming replication set 5-11
- replicate permissions 8-6
- replication on connection 8-3
- restoration 12-5

- resuming mirroring manually 7-7
- saving replication set 5-8, 5-9, 5-11
- selecting data for replication 5-8
- starting 2-2
- starting mirroring manually 7-7
- statistics 13-2
- stopping mirroring manually 7-7
- Throughput Diagnostics Utility (TDU) 6-8
- toolbar 2-5
- transmission
 - start options 10-7
 - stop options 10-8
 - window option 10-9
- transmission settings on connection 10-5
- verification log 9-9, 14-11
- verification on connection 9-7
- manual intervention** 11-27
- manual transmission** 10-2
- many-to-one configuration** 1-9
- map base** 6-2, 6-16
- map exact** 6-2, 6-16
- mapping**
 - all-to-one 6-2
 - custom 6-2, 6-16
 - Management Console 6-9
 - map base 6-2, 6-16
 - map exact 6-2, 6-16
 - one-to-one 6-2
 - text clients 6-16
- MaxLogFileSize** 13-32, 13-42
- MaxNumberOfLogFiles** 13-32, 13-42
- message window** 13-10
- mirror pause** 7-14, A-15
- mirror resume** 7-14, A-16
- mirror start** 7-13, A-16
- mirror stop** 7-13, A-17
- mirroring**
 - capacity 14-9, 14-34
 - how it works 7-5
 - Management Console
 - manually pausing 7-7
 - manually resuming 7-7
 - manually starting 7-7
 - manually stopping 7-7
 - on connection 7-6
 - on connection
 - Management Console 7-6
 - text clients 7-11
 - options 7-1
 - options compared 7-1, 9-3, 12-4
 - overview 1-3, 7-1
 - pausing 7-3
 - Management Console 7-7
 - text clients 7-14
 - resuming 7-3
 - Management Console 7-7
 - text clients 7-14
 - starting 7-3
 - Management Console 7-7
 - text clients 7-13
 - stopping 7-3
 - Management Console 7-7
 - text clients 7-13
 - text clients
 - manually pausing 7-14
 - manually resuming 7-14
 - manually starting 7-13
 - manually stopping 7-13
 - on connection 7-11
- modifying replication set**
 - text client 5-16
- monitor access** 15-1
 - Management Console 2-6
 - text clients 3-4
- monitor clear** 11-36, A-18

- monitor create** 11-28, A-18
- monitor delete** 11-36, A-18
- monitor display** 11-34, A-19
- monitor get** A-19
- monitor list** A-19
- monitor move** 11-29, A-20
- monitor option** 11-31, A-21
- monitor remove** 11-30, A-22
- monitor script add** 11-33, A-22
- monitor script remove** 11-34, A-23
- monitor start** 11-35, A-23
- monitor stop** 11-35, A-24
- monitor use** 11-29, A-24
- monitoring**
 - connection
 - Management Console 6-10
 - failover 4-4, 11-21
 - mirroring
 - Management Console 7-8
 - replication
 - Management Console 8-5
 - verification
 - Management Console 9-10
 - text clients 9-14
- monitoring tools**
 - DStat 13-30, 13-34
 - command prompt 13-34
 - log 14-11
 - Management Console 13-31, 14-4
 - Text Client 13-32
 - Event Viewer 13-14
 - LogViewer
 - log file 13-40
 - Management Console 13-41
 - text clients 13-42
 - viewing log files 13-43
 - Management Console
 - at-a-glance monitoring 13-6
 - display colors 13-9
 - message window 13-10
 - overview 13-2
 - statistics 13-2
 - overview 13-1
 - Performance Monitor 13-27
 - SNMP 13-48
 - Text Client
 - main window 13-12
 - text clients
 - status command 13-13
- mount point** 5-6, 6-2
- MSCS** C-1
 - see cluster
- multi-client access** 5-5

N

- niclist** 11-28, A-24
- no access** 15-1
 - Management Console 2-6
 - text clients 3-4

O

- one-to-many configuration** 1-10
- one-to-one**
 - active/active configuration 1-8
 - active/standby configuration 1-7
 - mapping 6-2
- orphan files** 7-3
 - Management Console 7-9
 - text clients 7-16
- orphans start** A-25
- orphans stop** A-25

P

- pause target** A-26
- Performance Monitor** 13-27
- permissions**
 - Management Console 8-6, 14-7
 - text clients 8-12
- ping** A-26
- placeholders** 11-25
- ports**
 - Failover Control Center 4-6, 11-26
 - Management Console 2-8
 - setting
 - Failover Control Center 4-6, 11-26
 - text clients 3-7
- program settings**
 - Management Console
 - database 14-10
 - general 14-2
 - logging 14-11
 - network 14-5
 - notification 14-12
 - queuing 14-6
 - source 14-7
 - startup 14-4
 - target 14-9
 - settings table for text clients 14-15, 14-37
 - text clients 14-14

Q

- queuing** 6-6
 - Management Console 6-13
 - text clients 6-23
- quit** A-28

R

- real-time transmission** 10-2
- recursive** 5-2, 5-9, 5-11, 5-13, 5-17, F-1
- remirror** 6-15, 6-24, 7-10, 7-18, F-1
 - Management Console 14-4
- renaming replication set** 5-11
- ReplicateNTSecurityByName** 8-12
- replication**
 - capabilities 5-6
 - how it works 8-1
 - Management Console
 - manually starting 8-4
 - manually stopping 8-4
 - monitoring 8-5
 - on connection 8-3
 - permissions 8-6
 - monitoring
 - Management Console 8-5
 - on connection
 - Management Console 8-3
 - text clients 8-8
 - options 8-1
 - overview 1-4, 8-1
 - permissions
 - Management Console 8-6, 14-7
 - text clients 8-12
 - starting
 - Management Console 8-4
 - text clients 8-10
 - stopping
 - Management Console 8-4
 - text clients 8-10
 - text clients
 - manually starting 8-10
 - manually stopping 8-10
 - on connection 8-8
 - permissions 8-12
- replication set rules** 5-2, F-2
- replication sets** F-2
- calculate size
 - Management Console 5-10
 - text clients 5-14
- copying 5-3, 5-11
- creating 5-3
 - text clients 5-12
- deleting 5-3
 - Management Console 5-11
 - text clients 5-19
- direct rule entry 5-9
- displaying 5-3, 5-16
- exclude 5-2, 5-9, 5-11, 5-13, 5-17
- how they work 5-4
- include 5-2, 5-9, 5-11, 5-13, 5-17
- listing 5-3, 5-16, 5-19
- Management Console
 - calculate size 5-10
 - copying 5-11
 - deleting 5-11
 - new 5-7
 - renaming 5-11
 - saving 5-8, 5-9, 5-11
 - selecting data 5-8
- modifying 5-3
 - text clients 5-16
- multi-client access 5-5
- new
 - Management Console 5-7
 - text clients 5-12
- operations 5-3
- overview 5-1
- properties 5-9
- recursive 5-2, 5-9, 5-11, 5-13, 5-17
- renaming 5-3, 5-11
- resync 5-3, 5-18
- revert 5-3, 5-18
- rules 5-2
- saving 5-3
 - Management Console 5-8, 5-9, 5-11
 - text clients 5-13, 5-18
- selecting data
 - Management Console 5-8
 - text clients 5-13
- text clients
 - calculate size 5-14
 - creating 5-12
 - deleting 5-19
 - displaying 5-16
 - listing 5-16, 5-19
 - modifying 5-16
 - new 5-12
 - resync 5-18
 - saving 5-13, 5-18
 - selecting data 5-13
 - wild cards 5-9
- replication start** 8-10, A-28
- replication stop** 8-10, A-28
- repset calculate** A-29
- repset create** 5-12, A-29
- repset delete** 5-19, A-29
- repset display** 5-16, A-30
- repset list** 5-16, 5-19, 6-16, 7-11, 8-8, A-30
- repset resync** 5-18, A-30
- repset rule add** 5-13, 5-17, A-31
- repset rule remove** 5-13, 5-17, A-31
- repset save** 5-13, 5-18, A-32
- repset use** 5-12, 5-16, A-32
- resource utilization** 1-2
- restoration**
 - automating 12-9
 - conditionals 12-3
 - conditionals compared 12-4
 - how it works 12-2
 - Management Console 12-5
 - options 12-3
 - options compared 12-4

- overview 1-6, 12-1
- text clients 12-7
- restore** 12-8, A-33
- resume target** A-34
- resync replication set** 5-18
- route** 6-2, 6-8
- rules** 5-2

S

- saving replication set**
 - Management Console 5-8, 5-9, 5-11
 - text clients 5-13, 5-18
- schedule**
 - options 10-3
 - window option 10-4
- schedule clear** 10-11, A-34
- schedule disable** 10-11, A-34
- schedule enable** 10-11, A-34
- schedule end** 10-14, A-35
- schedule start** 10-13, A-35
- schedule window** 10-15, A-36
- scheduled transmission** 10-2
- scripting** A-44
- security**
 - access levels compared 15-1
 - Active Directory 15-7
 - administrator access 15-1
 - Management Console 2-6
 - text clients 3-4
 - monitor access 15-1
 - Management Console 2-6
 - text clients 3-4
 - no access 15-1
 - Management Console 2-6
 - text clients 3-4
 - overview 15-1
 - Windows
 - how it works 15-2
 - Windows 200x and NT
 - user account 15-4
- selecting data for replication**
 - Management Console 5-8
 - text clients 5-13
- session interval** 10-3
- set** A-36
- setlocal** A-37
- share updates** 11-26
- shutdown** A-37
- single machine configuration** 1-12
- SNMP** 13-48
- source** 6-16, 7-11, 8-8, 9-11, 12-7, A-37
- source/target functionality**
 - current session 3-6
- sparse file** 6-2
- statistics**
 - connection
 - Management Console 6-10
- statslog start** A-38
- statslog status** 13-33, A-38
- statslog stop** A-38
- status** A-39

T

- target** 11-28, A-39
- testcon** A-39
- Text Client**
 - display 3-2
 - DTStat 13-32
 - monitoring tools
 - main window 13-12
 - overview 3-2
- text clients** 6-16, 8-10, 9-13
 - advanced settings 14-14
 - settings table for text clients 14-15

- auto-reconnect 6-24
- auto-remirror 7-18
- bandwidth limitations 10-12
- calculate replication set size 5-14
- clear transmission settings 10-11
- compression 6-22
- configuring failover 11-28–11-36
- creating replication set 5-12
- deleting replication set 5-19
- disconnecting connection 6-19
- displaying replication set 5-16
- listing replication set 5-16, 5-19
- LogViewer 13-42
- manually initiating transmission 10-17
- manually initiating verification 9-11
- manually starting replication 8-10
- manually stopping replication 8-10
- mirroring on connection 7-11
- modifying replication set 5-16
- monitoring tools
 - status command 13-13
- monitoring verification 9-14
- new connection 6-16
- new replication set 5-12
- orphan files 7-16
- pausing mirroring manually 7-14
- permissions 8-12
- ports 3-7
- program settings 14-14
 - settings table for text clients 14-15
- queuing 6-23
- replication on connection 8-8
- restoration 12-7
- resuming mirroring manually 7-14
- resync replication set 5-18
- saving replication set 5-13, 5-18
- selecting data for replication 5-13
- source/target functionality 3-6
- starting mirroring manually 7-13
- status command 13-13
- stopping mirroring manually 7-13
- Throughput Diagnostics Utility (TDU) 6-18
- transmission
 - start options 10-13
 - stop options 10-14
 - window option 10-15
- transmission criteria 10-11
- verification log 9-13

Throughput Diagnostics Utility (TDU)

- Management Console 6-8
- overview 6-4
- text clients 6-18

time limit 10-3

timenow A-40

transmission

- bandwidth limitations 10-2
 - Management Console 10-6
- text clients 10-12
- before connection 10-11
- clear settings
 - Management Console 10-5
 - text clients 10-11
- disable schedule 10-11
- enable schedule 10-11
- Management Console
 - bandwidth limitations 10-6
 - clear settings 10-5
 - manually initiating 10-10
 - start options 10-7
 - stop options 10-8
 - window option 10-9
- manual 10-2
- manually initiating
 - Management Console 10-10
 - text clients 10-17

- on connection 10-5
- overview 10-1
- real-time 10-2
- schedule
 - window option 10-4
- scheduled 10-2
 - options 10-3
- start options
 - Management Console 10-7
 - text clients 10-13
- stop options
 - Management Console 10-8
 - text clients 10-14
- text clients
 - bandwidth limitations 10-12
 - clear settings 10-11
 - manually initiating 10-17
 - start options 10-13
 - stop options 10-14
 - window option 10-15
- window option
 - Management Console 10-9
 - text clients 10-15
- transmission pause** 10-17, A-40
- transmission resume** 10-17, A-40
- transmission session start** 10-3
- transmission start** 10-17, A-40
- transmission stop** 10-17, A-41

U

- unload source** A-41
- unload target** A-41

V

- variables** 6-4, 6-21, A-44
- verification**
 - how it works 9-6
 - log 9-13
 - Management Console 9-9, 14-11
 - overview 9-4
 - Management Console
 - log 9-9, 14-11
 - manually initiating 9-8
 - on connection 9-7
 - manually initiating
 - Management Console 9-8
 - text clients 9-11
 - monitoring
 - Management Console 9-10
 - text clients 9-14
 - on connection
 - Management Console 9-7
 - options 9-2
 - overview 9-1
 - remirror options 9-2
 - remirror options compared 9-3
 - text clients
 - log 9-13
 - manually initiating 9-11
- verify** 9-11, A-42
- VerifyLogAppend** 9-14
- VerifyLogLimit** 9-14
- VerifyLogName** 9-13

W

- wait** A-43
- waitonmirror** A-43
- waitonrestore** A-43
- waitontarget** A-44
- wild cards** 5-9
- window** 10-4
- write** A-44